AD-A052 952 MASSACHUSETTS INST OF TECH CAMBRIDGE ARTIFICIAL INTE--ETC F/G 9/2
QUALITATIVE KNOWLEDGE, CAUSAL REASONING, AND THE LOCALIZATION O--ETC(U)
NOV 76 A L BROWN N00014-70-A-0362-0003

UNCLASSIFIED AI-TR-362 NL

1 OF 3
ADA
052952

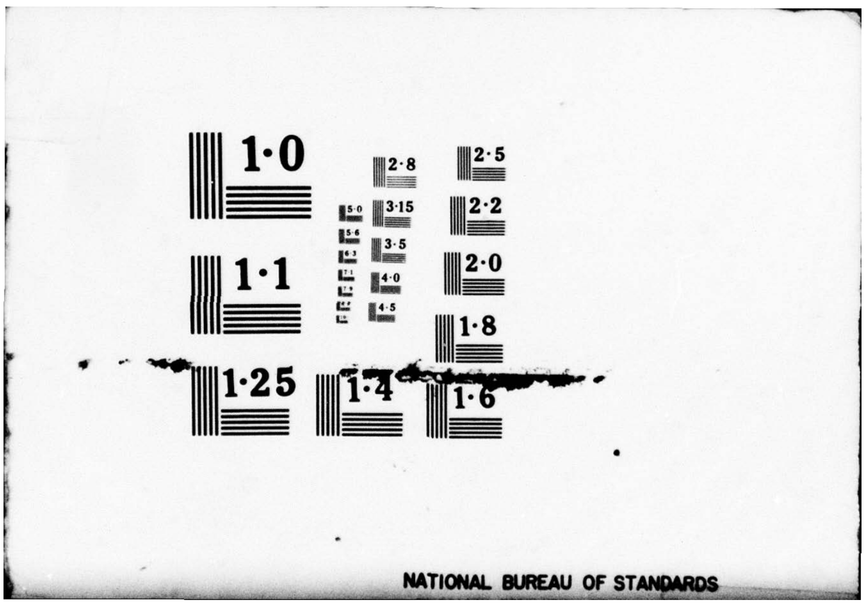
1·0
1·1
1·25
1·4
1·6
1·8
2·0
2·2
2·5
2·8
3·15
3·5
4·0
4·5
5·0
5·6
NATIONAL BUREAU OF STANDARDS

AI-TR-362

Qualitative Knowledge, Causal Reasoning, and the Localization of Failures

Allen Brown

~~March 1977~~

ARTIFICIAL INTELLIGENCE LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 14 AI-TR-362	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) 6 Qualitative Knowledge, Causal Reasoning, and the Localization of Failures.		5. TYPE OF REPORT & PERIOD COVERED Technical Report
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) 10 Allen Leon Brown, Jr		8. CONTRACT OR GRANT NUMBER(s) 15 N00014-70-A-0362-0003
9. PERFORMING ORGANIZATION NAME AND ADDRESS Artificial Intelligence Laboratory 545 Technology Square Cambridge, Massachusetts 02139		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Blvd Arlington, Virginia 22209		12. REPORT DATE 11 November 1976
		13. NUMBER OF PAGES 198
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research 12 193p. Information Systems Arlington, Virginia 22217		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Distribution of this document is unlimited.

9 Doctoral thesis,

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

None

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Automatic Maintenance
Electrical Analysis
Signal Tracing
Artificial Intelligence
Computer-edited Design

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

This report investigates some techniques appropriate to representing the knowledge necessary for understanding a class of electronic machines -- radio receivers. A computational performance model "WATSON" is presented. WATSON's task is to isolate failures in radio receivers whose principles of operation have been appropriately described in the knowledge base.

DD FORM 1473 1 JAN 73 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

UNCLASSIFIED 407 483

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

LB

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the Laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defence under Office of Naval Research contract N00014-70-A-0362-0003.

Qualitative Knowledge, Causal Reasoning, and the Localization of Failures

by

Allen Leon Brown, Jr.
Massachusetts Institute of Technology
November 1976

Revised version of a dissertation submitted to the Department of Electrical Engineering and Computer Science on September 3, 1975 in partial fulfillment of the requirements of the Degree of Doctor of Philosophy.

Abstract

This report investigates some techniques appropriate to representing the knowledge necessary for understanding a class of electronic machines -- radio receivers. A computational performance model - WATSON - is presented. WATSON's task is to isolate failures in radio receivers whose principles of operation have been appropriately described in his knowledge base. The thesis of the report is that hierarchically organized representational structures are essential to the understanding of complex mechanisms. Such structures lead not only to descriptions of machine operation at many levels of detail, but also offer a powerful means of organizing "specialist" knowledge for the repair of machines when they are broken.

Thesis Supervisor: Gerald J. Sussman

Title: Assistant Professor of Electrical Engineering and Computer Science

Acknowledgements

Many thanks are due Ira Goldstein, Marvin Minsky, Drew McDermott, and Gerald Sussman for their contributions to the technical content of this report.

O speculatore delle cose,
non ti laldare di conoscere le cose che ordinariamente
per sè medesima la natura conduce,
ma rallegrati di conoscere il fine di quelle cose
che son disegnate dalla mente tua.

– Leonardo da Vinci

Table of Contents

1 Introduction **1**

1.1 Why radios? 1
1.2 A theory of machine description. 6
1.3 Where do plans and designs come from? 8
1.4 A theory of machine diagnosis. 10
1.5 Related work. 13
1.6 What is to come. 17

2 Scenarios of WATSON's Performance **18**

2.1 A opened junction in a transistor. 20
2.2 A shorted capacitor. 29
2.3 A misaligned front end. 33
2.4 An opened collector-base junction. 36

3 A Preview of the WATSON Program **42**

3.1 Stating the problem. 42
3.2 Localizing the failure. 45
3.3 Describing radios to WATSON. 47
3.4 Programming constructs. 48

4 The Case of Q301 **51**

4.1 Lodging the complaint. 51
4.2 Signal tracing. 58
4.3 The loop problem solved. 64
4.4 The indictment of Q301. 69

5 The Case of C422 **73**

5.1 Getting started. 73
5.2 Localizing to the AGC. 75
5.3 Indicting C422. 79

6 A Case of Misalignment **83**

6.1 The usual preliminaries. 83
6.2 Localizing the problem. 84
6.3 Alignment. 87

7 The Case of Q407 **90**

7.1 Localizing to the audio driver amplifier. 90
7.2 Proposal and partial rationalization. 92
7.3 More DC rationalization. 97
7.4 AC rationalization. 99

8 Bug Abstraction. **100**

8.1 Q301. 100
8.2 C422. 106
8.3 The front-end. 112
8.4 Q407. 116
8.5 Using an abstraction. 116

9	**Signals**	**118**
9.1	Matching and propagation on the top-level plan-fragment	119
9.2	Propagation and matching in the RF section/AF section cascade	125
9.3	Propagation and matching in the RF section	129
9.4	Signals as voltages and currents	134
10	**Plans, Plan-fragments, and Designs**	**136**
10.1	Assembling the GR-78 design	136
10.2	Goals, purposes, and rules	139
10.3	Compiling determiners and agents	141
10.4	Needs and provisions	143
10.5	Further expansion of the RF portion of the GR-78 design	144
10.6	Compiling caveats	146
10.7	Plans with optional uses	149
10.8	A few words about components	151
10.9	Miscellaneous theorems	152
11	**Miscellaneous Details**	**154**
11.1	Scenario completeness	154
11.2	Consistency monitoring	156
11.3	LOOPS	157
11.4	Plan flattening	160
12	**Conclusions**	**162**
12.1	The underlying concepts	162
12.2	The machinery embodying the concepts	165
12.3	Some universals of problem-solving in deliberate systems	169
12.4	WATSON's shortcomings	175
12.5	Directions for further research	179
	Appendix: Index of Terms	**180**
	Bibliography	**183**

Table of Figures

1.1 The localization paradigm. 3
1.2 An overview of the recursive localization process. 5
1.3 A partial Winston network for a design. 9
1.4 Further interpretation of the recursive localization process. 12

2.1 The GR-78 circuit diagram. 19
2.2 An RF amplifier. 20
2.3 Flattened block diagram of the GR-78. 21
2.4 Top-level block diagram of the GR-78 with RF and AF sections. 22
2.5 An expanded RF section. 23
2.6 Converter expanded into an oscillator and mixer. 25
2.7 An expanded oscillator. 25
2.8 The oscillator circuit. 26
2.9 The DC subcircuit. 27
2.10 An expanded IF strip. 31
2.11 An expanded AGC2. 31
2.12 Circuit diagram of the low-pass filter. 32
2.13 An expanded RF amplifier. 35
2.14 An expanded AF section. 36
2.15 An expanded audio amplifier. 37
2.16 Circuit diagram underlying the audio amplifier. 38

3.1 Failure localization in WATSON. 44

4.1 The context structure at the time of validation. 55
4.2 A simple cascade. 62
4.3 A cascade with a simple part join. 62
4.4 A fork and join combined to give a loop. 63
4.5 LOCAL's process state. 65
4.6 Three elements forward and one back. 65

5.1 The circuit details of AGC2. 80

7.1 Expansion of the audio driver plan. 94
7.2 Double expansion of the audio power amplifier. 95
7.3 An expansion of Ampl. 96
7.4 An expansion of Amp2. 97

9.1 The ob hierarchy in the signal representation. 119
9.2 Basic control structure in SIG-MATCH. 121
9.3 A port commonly shared at two levels. 127
9.4 Flattening of the RF section modulo IFI. 132

12.1 The top-level plan of an automobile. 171
12.2 Expanded top-level plan. 172
12.3 An expanded power train 172
12.4 An expanded engine. 174

1 Introduction

What does it mean to understand how a machine works? How can such an understanding be used to diagnose and repair a machine when it is broken? In this report I shall be concerned with answering these questions in the context of diagnosis and repair of local failures -- failures whose ultimate etiology can be resolved to a single component, or distinguished collection of components -- in radio receiver circuits. A theory of what it means to understand a mechanism is realized in a design[1] for a program, called WATSON, which attempts to mimic the performance of a competent radio technician. WATSON is guided in the localization process by teleological and causal annotation associated with the design of a given radio receiver circuit.

1.1 Why radios?

The reader might well suggest that radio diagnosis seems quite simple; why study it? Is it not sufficient to have a table (possibly large) of underlying causes of failures, indexed by symptom[2]? The answer is negative for two reasons. In the first place, a different table would be required for each radio, hardly the acme of generality. The other serious defect is that the index

----

[1] WATSON is not presently a running program, though many of the features of his design have been incorporated in various programs at one time or another. This report will give detailed descriptions of how WATSON would go about isolating a number of non-trivial failures. The documentation of these successful "hand simulations" will presumably convince you that the design is realizable as a program.

[2] A symptom will turn out to be an association of a class of incorrect outputs with a class of correct inputs for a radio -- a 'class' being a generalization of a particular observation. Later when I reveal WATSON's internal representations for symptoms, you will see that indexing them is actually not so easy as it might seem.

is not unique -- different underlying causes lead to the same symptoms. As a practical matter, this latter defect underlies the observation that almost any noticeable failure in a radio results in a lack of audible output.

So things are not so simple after all. There are two possible patches to the basic table-look-up (TLU) strategy; but before we consider them, let's examine the paradigm that seems to underly any possible strategy for finding the bug. It is hypothesize and test (see figure 1.1). A failure localization strategy is an embedding of this paradigm in some control structure[1]. In the proposed TLU strategy, failure enumeration is limited by looking only at those failures known to have the observed symptoms. Testing consists of pulling the component (whose failure is alleged to be causing the problem) from the circuit and instrumenting it to see if it meets its intrinsic specifications. The non-uniqueness of the symptom index implies insufficient pruning of the search space of possible failures. Hence, on the average, many components will be pulled before the real culprit is discovered.

Another possible strategy, the "big crunch" (BC), is to make use of the physical theory of electronic machines embodied in Kirchoff's laws together with the voltage-current (VI) characteristics of individual components such as resistors, transistors, capacitors, and so forth. For each component on the radio's circuit diagram, a particular failure is proposed. This failure would change the VI characteristics of the component in a well-defined way. Testing the hypothesis would consist of solving the equations of motion for the circuit in the circumstances of the failure and matching the solution against the observed behavior of the circuit[2]. If the match succeeds the bug has been found. Whatever reservations one may have concerning the feasibility of making the necessary measurements, they are certainly overwhelmed by doubts about the feasibility of the necessary computations. A circuit of even moderate complexity would require an

----

1 Newell and Simon [1976] give an account in their Turing lecture of the importance of the hypothesize and test paradigm.

2 The observed behavior here is revealed by voltage-current measurements in all the various branches of the circuit.

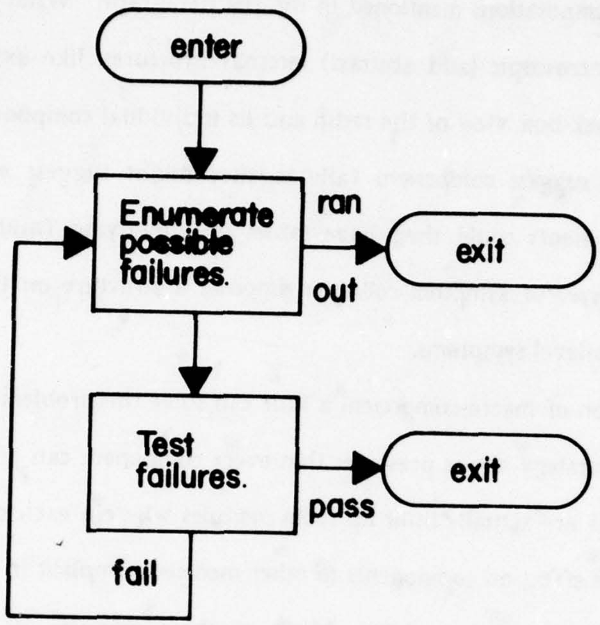

Figure 1.1 -- The localization paradigm.

astronomical amount of computation to carry out the BC strategy. Moreover, even when BC successfully uncovers the source of a bug, the answer is not very revealing. While BC will completely disambiguate failures that give rise to identical external symptoms, it provides so much detail that there is no way to formulate equivalence classes[1] of bugs.

Let's return to the problem of patching the TLU strategy. The problem is that when looking at the external symptomatology of the radio, we severely limit the available evidence. So it seems that we might ameliorate the problem of the many-to-one correspondence between bugs and symptoms by looking at some internal behavior of the radio as well. The question is, "What internal behavior?" The only distinguished sub-structures on the bare circuit diagram are components. Looking at symptomatologies with respect to such components implies doing

----

1 From a certain point of view, radios that differ at the circuit level will have identical abstract structures. Hence the different circuits will admit similar bug explanations modulo that viewpoint. This structural view gives rise to a partition of radio circuits according to the bugs they manifest.

precisely the same kinds of computations mentioned in the last paragraph. What seems necessary is the imposition of more macroscopic (and abstract) internal structures, like *amplifier*, *detector*, and *oscillator*, between the black-box view of the radio and its individual components. Instead of using external symptoms to suggest component failures, they might suggest *macro*-component failures. These macro-components could then have tables of underlying failures indexed by symptom. Adding another layer of symptom collection imposes a structure on the search space that can disambiguate the top-level symptoms.

Extending the notion of macro-component a little can solve the problem of generality as well. Consider that the BC strategy above presumes that every component can affect every other component. In reality designs are actually built up from modules wherein each component, more often than not, has negligible effect on components in other modules. Implicit in the radio circuit diagram is a hierarchical structure of modules. Many of these modules, like the amplifier, detector, and oscillator mentioned above, exist independently of a particular radio design This suggests making explicit the hierarchical structure implied by the design, and specifying the interactions among the modules mentioned in that structure. Hypothesis formation could then be driven by the causal and teleological commentary associated with the various modules. Testing would consist of matching the behavior suggested by the causal commentary against the behavior that is actually observed. I am suggesting the strategy depicted in the flow-chart of figure 1.2 where the interpretation of the steps is:

1. Does the macro-component, MC, meet its specifications?
2. Propose a part, P, of MC that might be broken.
3. Apply the localization process to P which may itself be a macro-component.
4. Exit with the name of the failing part found at step 2.
5. Exit, complaining of false accusations.
6. Exit, complaining that there is no part to take the blame.

This is a variant of BC that need never ask questions about the insides of macro-components

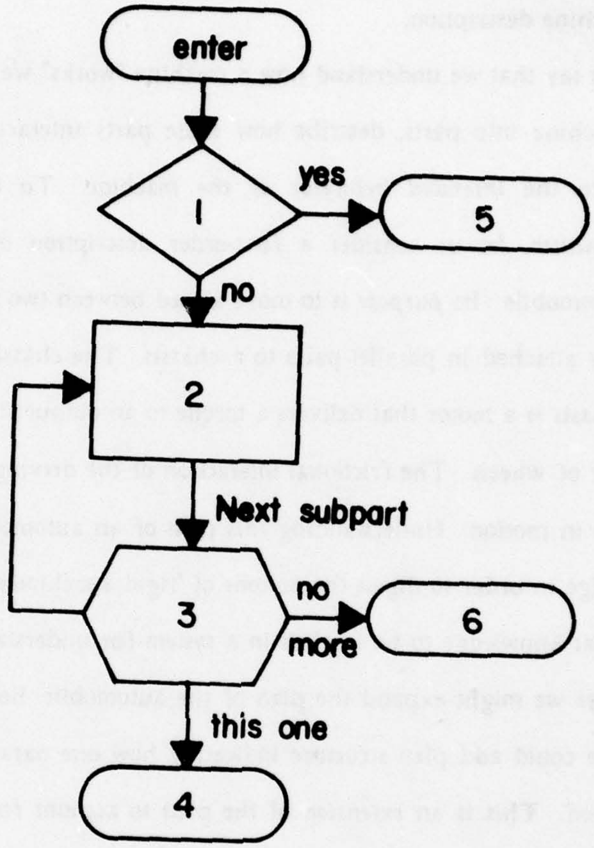

Figure 1.2 – An overview of the recursive localization process.

unless there is some good reason to suspect them of causing trouble. Moreover, hierarchical structuring of BC allows schemes for proposal of failures to be as specifically tailored as is necessary to the particular kind of macro-component under investigation.

Radio repair offers a reasonable micro-world for understanding how machines work, because a general (feasible) strategy for carrying out such repairs requires the understanding of the function of the whole in terms of the functions of some collection of parts. In particular, WATSON's success at finding bugs in radios depends on his comprehension of design descriptions together with techniques of causal reasoning which are driven by, and manipulate those descriptions.

1.2 A theory of machine description.

When we say that we understand how a machine "works" we usually mean that we can decompose the machine into parts, describe how those parts interact, and indicate how those interactions lead to the *intended* behavior of the machine. To motivate the problem of representing mechanism, let us consider a first-order description of a machine familiar to everyone -- the automobile. Its *purpose* is to move a load between two places. To that end it has four wheels rigidly attached in parallel pairs to a chassis. The chassis supports the load. Also attached to the chassis is a motor that delivers a torque to an output shaft. The shaft is coupled to one parallel pair of wheels. The frictional interaction of the driving wheels with the road sets the entire machine in motion. Understanding this plan of an automobile presumes a great deal of physics knowledge in order to digest the notions of 'rigid attachment' and 'rotational coupling'. We will assume that knowledge to be implicit in a system for understanding mechanical devices. There are two ways we might expand the plan of the automobile: horizontally or vertically. In the former case we could add plan structure indicating how one parallel pair of wheels enables the car to be steered. This is an *extension* of the plan to account for additional function. An instance of vertical expansion would be a sub-plan detailing the parts of the motor, indicating how these parts interact to yield a torque at the output shaft. This is a *refinement* of the plan to give a more detailed explanation of some function.

Let's try to make an analogous description for a radio receiver. The purpose of the receiver is to select one of several possible signals from the "ether," demodulate it and transduce the modulation into sound. To that end a radio has two principle parts, a radio frequency (RF) section and an audio frequency (AF) section. The output of the RF section is coupled to the input of the AF section via a signal port. The RF section selects the signal of interest, demodulates it, and delivers the naked modulation signal to its output port. This modulation is processed by the AF section via amplification and frequency response equalization (e.g. RIAA, NAB, and various tone controls) to produce a signal suitable for conversion to sound by the loudspeaker. As with

an understander of automobile plans, we again assume that an understander of radio circuit plans has an implicit comprehension of the atomic concepts. Hence signal, signal coupling, and signal port are notions embedded in WATSON's procedures. A concept like "frequency selection" would be built up from more atomic concepts. Such a concept would be represented explicitly in the plan. As before, we can consider extenstions and refinements of the radio plan. We might add another part to the plan, a power supply. Or we might consider a sub-plan for the RF section containing an RF amplifier, a converter, an intermediate frequency (IF) amplifier, and a detector, with a suitable explanation of how they interact to yield the overall function of selectivity and demodulation.

WATSON must be told more about a broken radio than its observed symptoms. To begin with he wants a complete circuit diagram of the radio. This diagram must then be *mapped* into a space of plans (which are types); that is, he needs its design. I call this mapping process *binding*. At each level in the binding process, parts are bound in the context[1] of bindings at higher levels. A top-level plan-fragment -- a token of a plan (perhaps like the one described in the last paragraph) -- is chosen. The radio is associated with this plan-fragment in an empty context. The plan, of which the plan-fragment is an instance, has a number of parts. Plan-fragments are chosen in turn for each of these parts. That is, each part is bound to its chosen plan-fragment in the context of the previous binding. This process is recursive. The product of this process is a tree of bindings. A branch of this tree is terminated by binding a part to a plan-fragment -- whose type might be the abstract resistor -- which cannot be decomposed further. A tree of bindings that has been *closed* in this manner is called a plan closure. Terminal elements in the tree may be associated with components on the radio's circuit diagram. The plan closure together with the component association and ancillary commentary is called the *design* of the radio. Plans

----

1 A part may be viewed as a kind of variable which takes on a value in a design. Because the design induces a hierarchical structure on parts, the definitions of parts near the leaves of the hierarchy may freely refer to definitions near the root of the hierarchy, hence the the usual scoping conventions of block-structured programming languages comes to mind.

distinguish interfaces of part interactions (such as ports and terminals) and the modality of interaction (such as signal and branch current). The input-output behavior of the parts, as viewed from their interfaces, is represented by an implicit plan sub-structure called an *io-contour*. The io-contour also indicates what overall effect results from the interactions of the plan's parts[1]. A plan may carry annotations concerning the external requirements to be met (or circumstances avoided) in order for it to perform its specified task. Finally, plans serve as the structures in terms of which failure mechanisms are abstracted. Keep in mind that since designs are built up from plan-fragments, the attributes of the types of the latter are inherited by the radios in whose designs their instantiations are bound. Figure 1.3 gives an illustrative relational network [Winston, 1975] of the kind that represents a design. INSTANCE-OF relations show how tokens of plans are given. A combination of PART-OF and INSTANTIATES-PART-IN relations yields a binding. Finally, the TOP-LEVEL-PLAN-FOR relation gives the name of a design and a plan-fragment used for its top-level description. One of the many descriptive features of designs not shown in this network is the fact that plan-fragments (the circular shapes in figure 1.3) can be identified with components on the circuit diagram.

1.3 Where do plans and designs come from?

Why must WATSON be told the design of the radio? Doesn't a good technician infer the design (from the circuit diagram) for himself? The circuit diagram of a radio typically supplies the technician with many hints about the radio's modular structure. Radio circuit diagrams have labels on them like "RF amp," "IF section," "IF can," "detector," etc. These labels denote goals which may be achieved by plans that are well known in the culture of radio design. The layouts of the circuit diagrams themselves are stylized so as to give certain kinds of information about the workings of the circuit, e.g. DC voltages (especially biases) tend to "fall" from top to bottom.

---

1 The io-contour may be thought of as a distinguished sequence of verification conditions of the variety that appear in the Floyd-Hoare-style axiomatic semantics for programming languages.

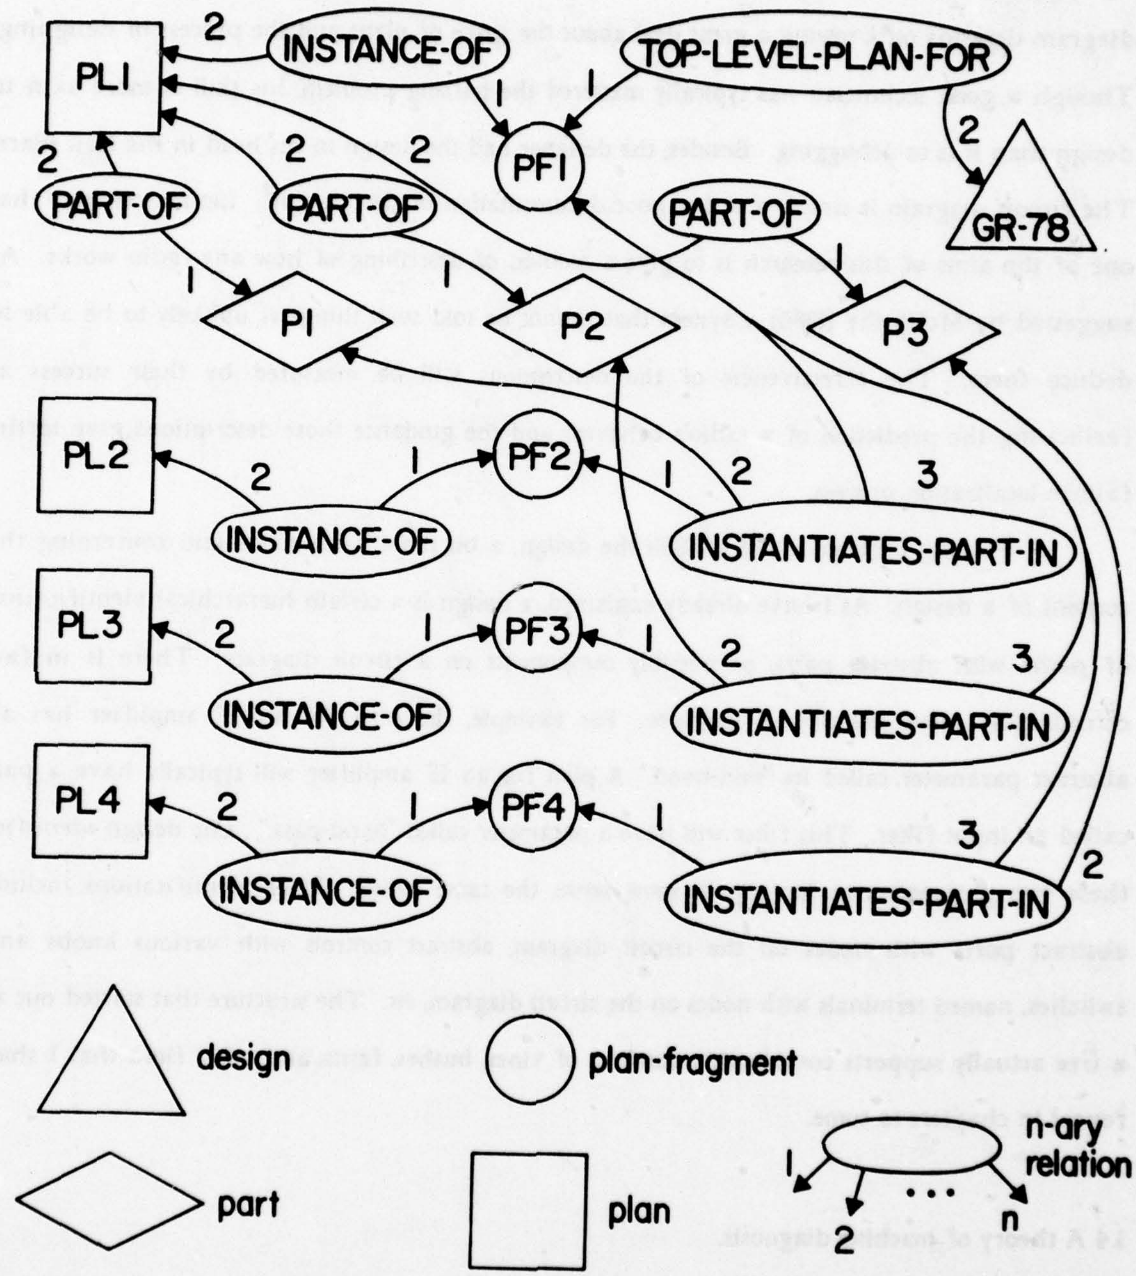

Figure 1.3 -- A partial Winston network for a design.

and signals "flow" from left to right. The parsing process that yields a design from a bare circuit diagram depends on knowing a great deal about the space of plans and the process of designing. Though a good technician has typically mastered the parsing problem, his skill is more akin to design than it is to debugging. Besides, the designer had the design in his head in the first place. The circuit diagram is simply a rather poor documentation of that design[1]. Let me reiterate that one of the aims of this research is to give a method of describing of how any radio works. As suggested by McCarthy [1968], a system that cannot be told such things is unlikely to be able to deduce them. The effectiveness of the descriptions will be measured by their success at facilitating the prediction of a radio's behavior and the guidance those descriptions give to the failure localization process.

Having decided to tell WATSON the design, a bit more needs to be said concerning the content of a design. As I have already explained, a design is a certain hierarchical identification of plans with abstract parts, or possibly components on a circuit diagram. There is in fact considerably more identifying structure. For example, the plan for an IF amplifier has an abstract parameter called its "mid-band." A plan for an IF amplifier will typically have a part called an input filter. This filter will have a parameter called "band-pass." The design identifies these two parameters as having, in some sense, the same value. Other identifications include abstract ports with nodes on the circuit diagram, abstract controls with various knobs and switches, named terminals with nodes on the circuit diagram, etc. The structure that started out as a tree actually supports considerable numbers of vines, bushes, ferns, and other flora that I shall reveal in chapters to come.

1.4 A theory of machine diagnosis.

---

1 One of the reasons for the poor documentation is the lack of a good formal language in which to do the documentation. WATSON proposes, among other things, a solution to this documentation problem.

The localization of failures, apart from being an interesting problem in its own right, serves as a test of whether WATSON has truly understood the machines that have been explained to him. WATSON's methods of failure localization, as suggested earlier, are applied recursively to the plan-fragments bound in the design of the failing radio. The flow-chart of figure 1.4 shows the recursive control structure of the failure localizing process in greater detail. Starting with the plan-fragment (referring to figure 1.3 may be helpful) bound at the top-level to the radio, he checks his notes to see if he has previously abstracted any bugs (for that plan-fragment's type) whose symptom would match the observed symptom[1]. The abstracted bug, among other things, associates a symptom for the plan, a part of the plan, and a *sign* for that part. The sign is a symptom for the part. In order to put the blame on that part, the part must exhibit the correct sign. If the sign is present, the localization is recursively activated by dredging up the plan to which that part is bound, with the sign becoming the symptom at the new level of recursion.

If the indicated sign is not present, there may be other bugs known to WATSON whose patterns of symptomatic behavior make them applicable in the present situation. WATSON exhausts all such candidates before proceeding to other hypothesis-generation strategies. One such strategy, LS2 , relies on the essentially causal nature of the machines under consideration. Plans may give the structure of the *flow* processes being carried out in the machines they explain. Such processes can be *back-traced*[2]. That is, starting at the final output, and knowing what process each part is designed to accomplish, WATSON can work his way along the processing path(s) until he finds a part whose output is incorrect given its input(s). If radio plans were strictly causal, back-tracing would always lead directly to localizing the failure. Unfortunately, plans have parts which are not really *unilateral*[3]. Moreover, the flow paths in plans may have

- - - - - - - - - -

[1] Looking to see if the answer is already known is one of several possible failure localization methods or strategies. I shall refer to it as the "LS1" strategy.

[2] Note that implicit in the back-tracing process is the assumption of (among other things) the independence of parts in different modules, about which I remarked earlier.

[3] A part in a flow process is unilateral if its output side cannot affect its input side. Indeed, excluded from this class are various passive filter networks.

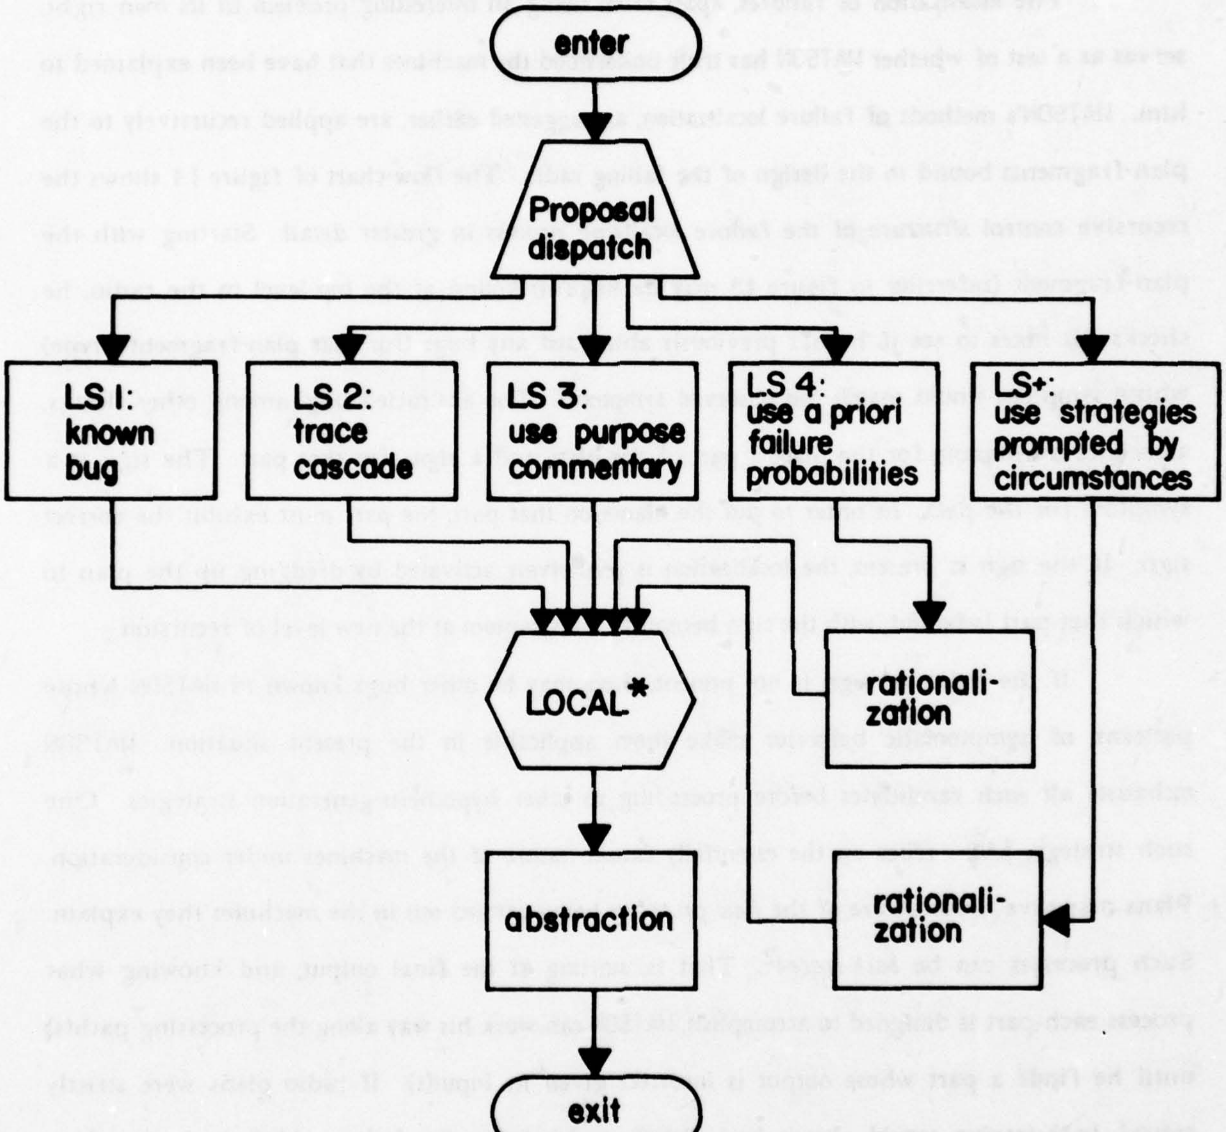

Figure 1.4 -- Further interpretation of the recursive localization process.

feedback. Such non-idealities necessitate other hypothesis-generation mechanisms.

The next method, LS3 , works by considering the purposes of various parts in a plan.

A *purpose* is represented by a predicate whose truth is guaranteed by a *rule*. Networks of these rules and predicates make up the io-contour mentioned earlier. The io-contour, then, constitutes a kind of electronic calculus that plays a role at the module level analogous to that played by differential equations of motion at the component level. Forming the hypothesis consists of proposing that some predicate is no longer true (i.e. its supporting rule has been invalidated) and predicting the consequences of such a turn of events. Another hypothesis-generation mechanism, LS4 , considers the circuit components associated with the terminal plan parts and suggests failures in those components based on the *a priori* probabilities of the components' succumbing to known faults. LS3 and LS4 have the property that the hypotheses they produce must be *rationalized*. Unlike back-tracing and hypothesizing previously abstracted bugs, these other mechanisms can hypothesize failures that do not necessarily lead to the observed symptom. Hence the consequences of the hypothesis must be deduced and shown to agree with the observed symptom. When the hypothesis-generating mechanisms succeed in producing a plausibly failing part, WATSON abstracts the failure mechanism (if it is previously unknown) and invokes the localization process recursively. The recursion terminates at such time as a localized part is found to be associated with a circuit component and *verified* by removing the component and examining it for the proposed failure. It goes without saying that at any level of localization contradictory evidence may be discovered that necessitates backing out of an hypothesis[1].

1.5 Related work.

There are a number of programs whose motivations or methodologies are close enough to WATSON's to warrant comparison and contrast. They fall into two broad categories: programs whose principal concerns are debugging, diagnosis or "linear" problem-solving; and programs

---

[1] There are also localization strategies that are specific to particular plans. Such strategies will be referred to collectively as LS+. The notion here is that there tricks that frequently work in radio repair (like inspecting the heater filaments in tube radios) that are not sufficiently universal to be subsumed by bug abstraction.

that are essentially computer-aided analysis systems charged with yielding predictions about the dynamic and/or static behavior of particular electronic circuits. The first category includes Sussman's HACKER [1975], Goldstein's MYCROFT [1974], Sacerdoti's NOAH [1975], Shortliffe's MYCIN [1974 and 1976], Brown's SOPHIE [Brown et al., 1974, Brown et al., 1975], and de Kleer's trouble-shooting program [1975]. The second category includes Penfield's MARTHA [1971], Dertouzos' CIRCAL [1967], and EL, a recent effort of Sussman and Stallman [1975], and Roylance [1975]. WATSON's immediate ancestors are HACKER and MYCROFT. HACKER is a program that becomes increasingly skilled at solving blocks-stacking problems. He does so by debugging old programs to fit new situations. WATSON, in contrast, becomes increasingly skilled at debugging radios. He does so by enlarging his repertoire of bug classifications. Since the bug classifications are abstracted in terms of plans which may be bound in many radio designs, the new knowledge is available in any of those designs. HACKER tries to explain bugs in terms of pre-compiled classes. The success of both programs is attributable to the extensive commentary on the machines of their respective domains. This commentary indicates what the parts are trying to accomplish and for whom. MYCROFT is an expert at debugging a limited class of LOGO programs. The class of programs in question has a particularly structured product -- pictures. Goldstein shows how models of the intended product, together with the user written program, can be used to infer the user's plan for his program. Since bugs generally arise at the interfaces of plan steps, knowledge of potential bad interactions among those steps allows MYCROFT to repair the plan, and thence the program. He can also repair some programs that are not correct implementations of a plan. WATSON similarly relies on the causal and teleological commentary embedded in plans to guide the debugging process. His job, however, is to convert a perfect plan into one that would exhibit the observed symptoms. The fundamental advance of WATSON over HACKER and MYCROFT is two-fold. In the first place, WATSON has a qualitative theory of how the machines of interest to him work. He can make predictions about their behavior by means other than running them. Or more accurately, he models machine behavior at various levels of detail. In the second place, he

exercises hierarchical control of local reasoning. This keeps him out of the trouble that theorem provers classically encounter. No more facts pop into his head than are relevant at the time. As it turns out, the phenomenon of locality will be the key ingredient in both the representational mechanisms and control structures embodied by WATSON.

NOAH is a producer of plans while WATSON is mostly a user (though sometimes incremental modifier) of plans. Despite their different goals, they both rely on hierarchical plan representations. Sacerdoti's procedural net has many of the same attributes as my space of plans. His nodes (my plans) represent various levels of detail in actions (processing) used to modify the environment. Both systems apply various evaluators to substructures of their plans to compute what effect a plan has on its environment. The differences in our plan structures seem to arise mostly from his interest in manipulating the *temporal* structure of plans in contrast to my interest in manipulating their *action* structure.

Shortliffe's system exhibits a number of design choices similar to those I have made in WATSON. My symbolic description of signals is not unlike his parameterized descriptions of patients. His system relies on a system of rules to deduce correct therapies. I have a system of rules for matching signals and branch variables (currents and voltages), and a system of rules for propagating signals and branch variables to do causal reasoning. More importantly, MYCIN's and WATSON's uses of rules differ in that in the former system, the rules are the diagnosis whereas in the latter rules are used as predictive aids in guiding the diagnostic process. WATSON attempts to embody a general class of diagnostic techniques to be applied to particular machine descriptions, whereas MYCIN makes use of a collection of diagnostic techniqes directed specifically at the artifact being debugged. It should also be reiterated that WATSON's hierarchical representational scheme allows rules to be understood in terms of (presumably) more primitive rules. Essentially MYCIN has no theory other than the rules of diagnosis. If a physiological theory were added, the resulting system would be akin to WATSON.

Since SOPHIE is a system for the teaching of trouble-shooting in electronic instruments,

it is reasonable to expect that she must confront many of the same technical issues as confronted by WATSON. Her very different solutions to the problems seem to stem from a number of factors: desire for speed, general question answering ability, concentration on one particular instrument, and the fact that that instrument is a power supply[1]. SOPHIE's deductive methodology seems to rely on converting essentially qualitative questions to quantitative ones that can be posed to a simulator, augmented with specific knowledge of a particular power supply. The representational machinery is geared to supporting that interface and to lending semantic support to a natural language front end.

De Kleer is concerned, as I am, with keeping the trouble-shooting program's reasoning confined to locally available knowledge. He restricts his qualitative methods to DC circuitry, however. This is partly because AC qualitative analysis seems to require hierarchical, teleological structures (which he does not investigate), and because he is interested in optimal measurement strategies and the purposes that underlie measurements.

Sussman, Stallman, and Roylance have recently reported on a new circuit analysis program called EL. EL makes use of causal reasoning in much the same local fashion as does WATSON. Since EL's plan-fragments have no "insides," EL cannot reap the benefits of hierarchical reasoning. This means, of course, that EL's notion of causality is "flat," a circumstance that limits the complexity of the mechanisms to be analyzed as well as the depth of their analysis.

MARTHA and CIRCAL are classical analytic aids to circuit design. The goals underlying their designs are quite different from those underlying WATSON's. They embody a great deal of knowledge about analytic models of electronic circuits. While WATSON finds it straightforward to summarize a radio's qualitative behavior at various levels of detail (a synthetic problem), MARTHA's and CIRCAL's major stumbling blocks arise when trying to extract what they know in terms of the designer's problem. Conversely WATSON cannot explain a fourteen-pole Butterworth filter

----

1 A power supply is a machine that is in a sense designed to exhibit no dynamic behavior. Hence its behavior is easily characterized by numbers when operating correctly.

(because he has no theory of interaction for the filter's parts), while the analytic systems find detailed mathematical explanations (Remember the BC strategy!) of such devices straightforward Clearly what is needed is an engineer's aid that is a synthesis of the two kinds of systems.

1.6 What is to come.

The remainder of this report hinges heavily on the next chapter. It gives an informal account of how WATSON goes about his business. Beyond that, chapters are organized more or less breadth-first. I attempt to give appropriate forward pointers so that you may examine any aspect of the design in as much detail as suits your fancy ...

2 Scenarios of WATSON's Performance

In this chapter I shall present four scenarios of WATSON's execution of the trouble-shooting task. These scenarios will be informal in two senses. First, nothing will be said for now about how we might communicate to WATSON the nature of the faulty radio's symptoms or the radio's design. The second sense of informality lies in leaving unrevealed for the time being the precise representations of WATSON's knowledge and deductive methods. The scenarios will be presented as if they were protocols of a human technician. The purpose of this chapter is, therefore, to acquaint you with what WATSON can do, rather than how he does it.

Each scenario is concerned with a local failure in a particular radio receiver, the Heathkit® GR-78 [Heath Company, 1969] whose circuit diagram is presented in figure 2.1. Although the receiver actually covers six bands between 190 kHz and 30 MHz, and is capable of amplitude modulated (AM), continuous wave (CW), or single sideband (SSB) reception, WATSON's model of the radio is that it is an AM receiver for the band between 3 MHz and 7.5 MHz. This simplification entails no loss of generality since WATSON already has mechanisms for representing variability in the configuration of the radio (see chapter 10). These mechanisms could readily handle a more complete model of the GR-78, including switching bands and changing the mode of demodulation. Moreover, the representations and methodologies do not depend on a particular receiver. The last observation addresses a more important question that might be raised concerning the use of a single instrument as a source of examples. Might not this indicate the WATSON's methods are limited to that instrument? The GR-78 was chosen in the first place because it is a receiver of moderate complexity. As such it incorporates in its design a substantial set of plans from the space of possible radio circuit plans. Also the methods employed by WATSON have been hand-simulated in the trouble-shooting of circuits other than the GR-78.

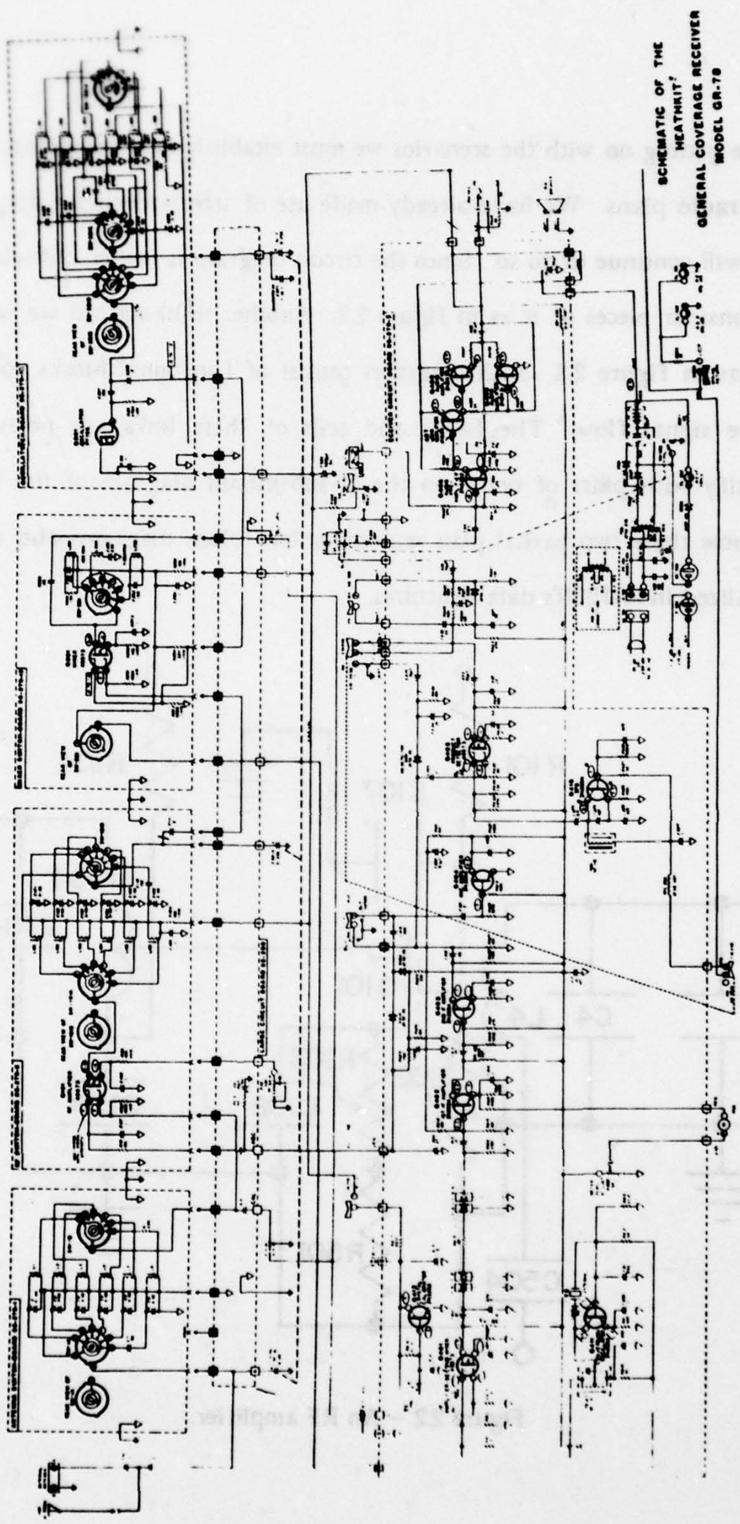

Figure 2.1 -- The GR-78 circuit diagram.

Before getting on with the scenarios we must establish some temporary conventions for thinking about radio plans. We have already made use of standard circuit diagram notation in figure 2.1. We will continue to do so. Since the circuit diagram is rather unwieldy in its entirety, we will often consider pieces of it as in figure 2.2. Another thinking aid we will employ is the block diagram as in figure 2.3. Such diagrams consist of functional blocks connected by links representing the signal flow. The heads and tails of these links are ports which we will sometimes identify with pairs of nodes on the (a sub-)circuit diagram of the GR-78. We will eventually see how these two partial plan representations (block diagrams and circuit diagrams) are formally realized in WATSON's data structures.

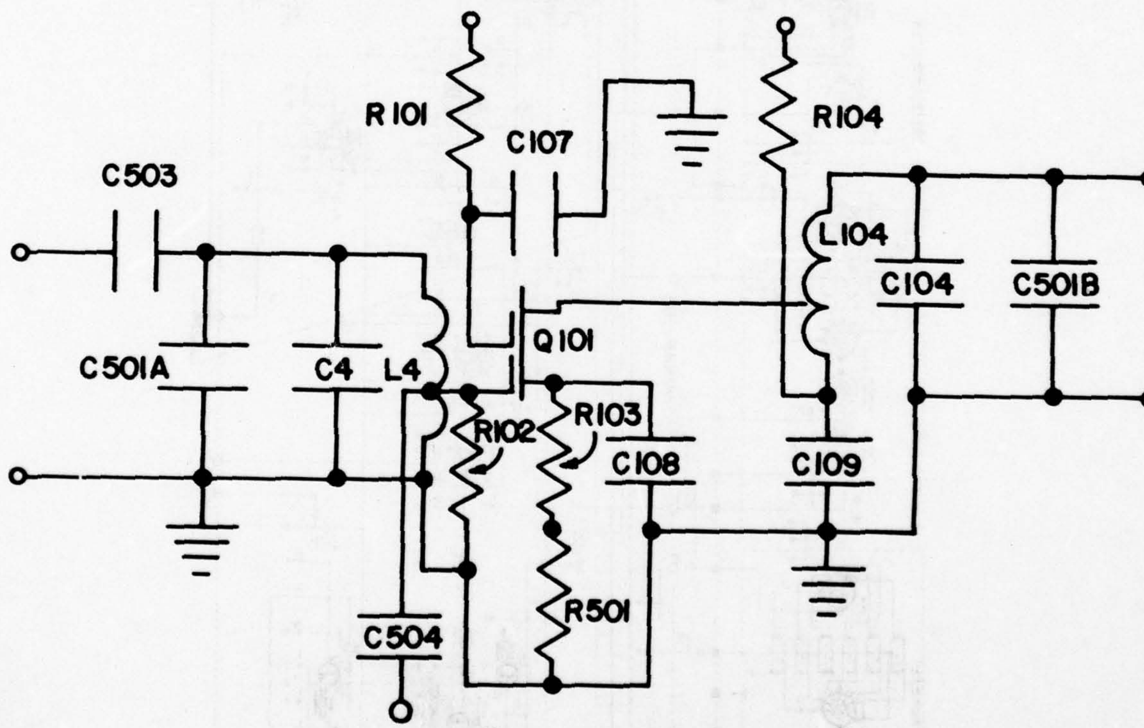

Figure 2.2 -- An RF amplifier.

2.1 A opened junction in a transistor.

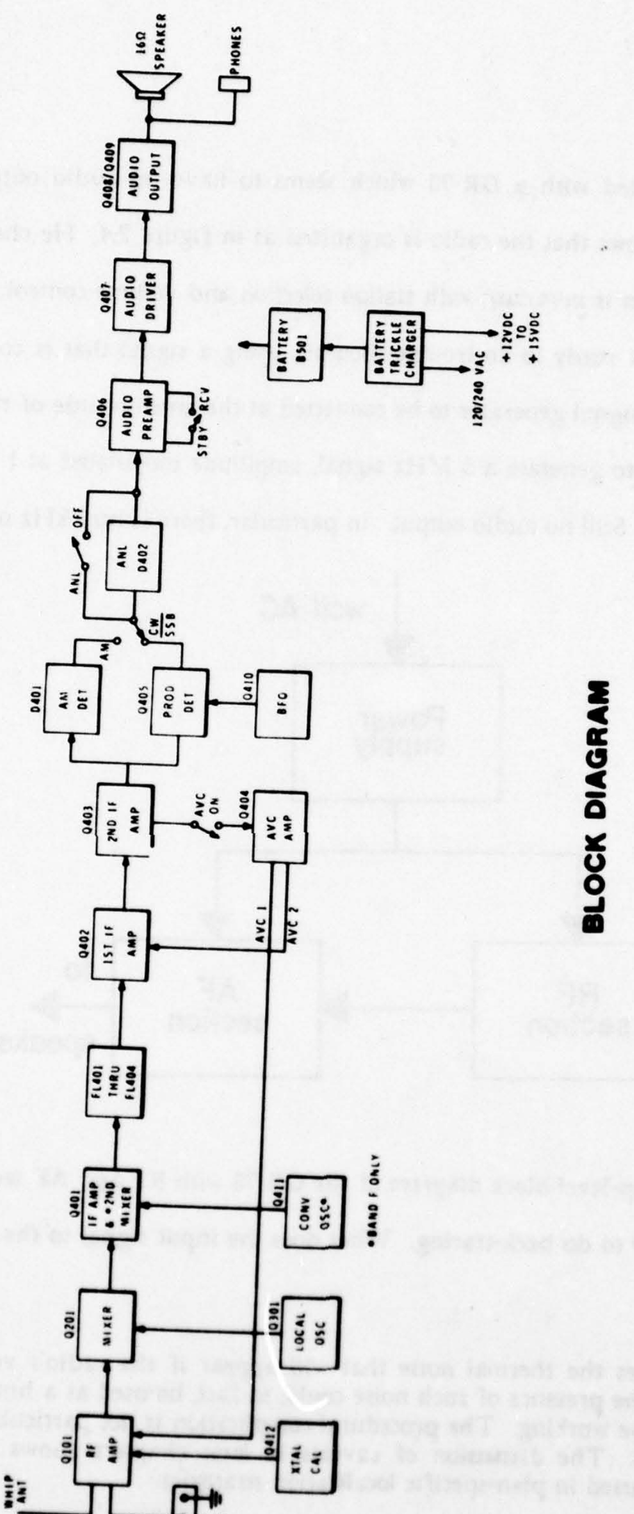

Figure 2.3 -- Flattened block diagram of the GR-78.

WATSON is presented with a GR-78 which seems to have no audio output under any circumstances[1]. WATSON knows that the radio is organized as in figure 2.4. He checks to see that the output of the AF section is invariant with station selection and volume-control setting. This is indeed the case. He is now ready to do trouble-shooting using a signal that is completely under his control. He calls for a signal generator to be connected at the antenna side of the RF section[2], and requests that it be set to generate a 5 MHz signal, amplitude modulated at 1 kHz and tunes the GR-78 around 5 MHz. Still no audio output. In particular, there is no 1 kHz output signal.

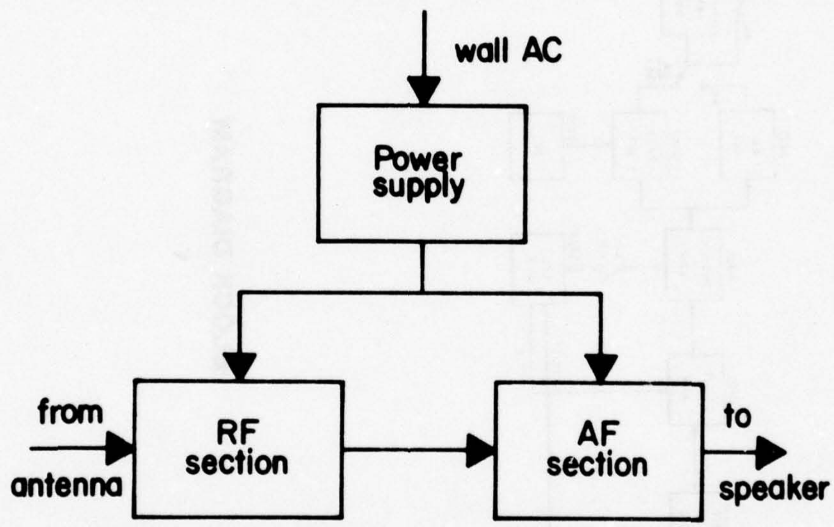

Figure 2.4 -- Top-level block diagram of the GR-78 with RF and AF sections.

Now he is ready to do back-tracing. What does the input signal to the AF section look

----

1 This observation ignores the thermal noise that will appear if the radio's volume control is turned up sufficiently. The presence of such noise could, in fact, be used as a hint by WATSON the power supply is likely to be working. The procedural complication is not particularly revealing in the present circumstances. The discussion of caveats in later chapters shows how such hints might be formalized and used in plan-specific localization strategies.

2 Notice that I have not mentioned The AF section's input from the power supply -- a power port. In general I shall not mention back-tracing through power ports (unless they are the source of a problem), as checking them is monotonously straightforward. Assume, however, that in every recursive application of the localization process such ports are in fact being checked.

like? WATSON knows that there should be a 1 kHz signal there, but there is not. So the trouble is probably not in the AF section, and therefore must (discounting the possibility of a shorted input port[1]) originate somewhere before it. Hence the RF section should be examined next. WATSON knows that the RF section's input signal must be correct since the signal generator is supplying it. The trouble then is most probably in the RF section. WATSON must now examine the organization of the RF section, which is presented in figure 2.5.

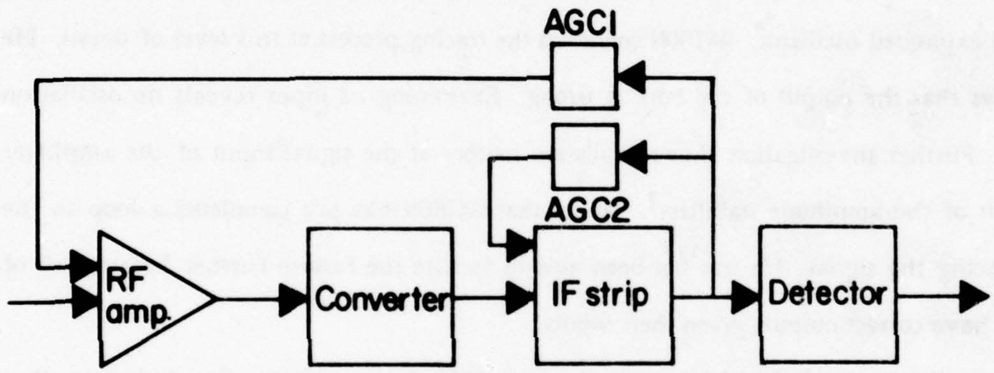

Figure 2.5 -- An expanded RF section.

Since the output port of the detector is the same as the output port of the RF section, the signal at that port is already known to be bad (i.e. not a 1 kHz signal). Looking at the input port of the detector, WATSON discovers that it is not a 455 kHz signal amplitude modulated at 1 kHz, as the operational theory of this radio predicts that it should be. The preceding module is the IF section, whose functional semantics demand that its input too should be a 455 kHz AM signal (and perhaps some spurs), though having a smaller amplitude than its corresponding output. Theory does not correspond to reality in this case. Back-tracing again, WATSON arrives at the converter. The theory of the GR-78 indicates that the input signal should consist mostly of a 5 MHz component, amplitude modulated at 1 kHz. Inspection verifies this[2].

- - - - - - - - -

1 This possibility is "pushed." If the possibility of "before AF section" fails to pan out for some reason, WATSON will return to this point.

2 Note that this last measurement is difficult to do without the aid of a very elegant RF measuring instrument (which turns out to be another receiver in essence). The control structure employed by most technicians would probably hypothesize that the input to the converter was, in

WATSON now has his hands on a module whose inputs seem correct, but whose output is not. It is a good bet that the trouble is inside -- which calls for expanding the inside. Figure 2.6 shows the main ingredients of a converter as seen by WATSON. The mixer has two signal inputs. WATSON has already verified that one of these inputs is correct. The other should be an unmodulated sinusoid of frequency 455 kHz plus the frequency of the broadcast signal -- 5.455 MHz in this case. It is not. Since the oscillator has no inputs, it is very likely the culprit. Figure 2.7 shows an expanded oscillator. WATSON continues the tracing process at this level of detail. He already knows that the output of the tank is wrong. Examining its input reveals no oscillation there either. Further investigation shows oscillation neither at the signal input of the amplifier nor the input of the amplitude stabilizer[1]. Notice that WATSON has just completed a loop in the course of tracing the signal. He has not been able to localize the failure further because all of the modules have correct outputs given their inputs.

This situation calls for the invocation of WATSON's loop specialist. One technique that can be used in such situations is to break the loop and independently supply the signal that should be present at the break point. Ordinary tracing can then be used to find the module that is causing the problem. A variant on this technique involves the injection of a signal at some point in the loop without making a physical break. WATSON can then check to see whether the signal propagates forward correctly. Loop breaking techniques are the loop specialist's *forte*, but they are expensive to use. So WATSON first tries for a "quick and dirty" solution using an *ad hoc* technique.

----

fact, okay and continue. The holding of some hypotheses in abeyance while pursuing other hypotheses in parallel is easily achieved using CONNIVER's generalized control structure. For the purposes of the present exposition, we shall presume that WATSON will always choose to motivate hypotheses by measurements rather than optimistically assuming that a hypothesis will work out in the absence of such motivation.

[1] The sense of 'input' here is purely abstract. Standing alone, the RC network underlying the amplitude stabilizer has no input or output. Viewed as a two-port network it is quite symmetric. Its *use* in the oscillator suggests an input side and an output side.

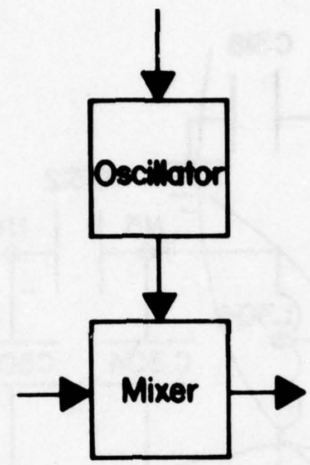

Figure 2.6 -- Converter expanded into an oscillator and mixer.

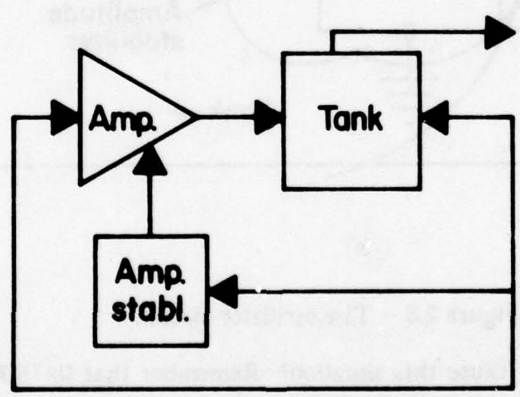

Figure 2.7 -- An expanded oscillator.

The heuristic used is to consider each of the active modules[1] in the block diagram. These modules get expanded to the level of detail of the circuit diagram where the consistency of the bias voltages may be examined. A quick check in the present case reveals that the gate bias of the junction field effect transistor (JFET) Q301 in figure 2.8 is is inconsistent with the source

---

1 An active module is one with power ports in addition to signal ports. The amplifier in figure 2.7 actually has a power port associated with it.

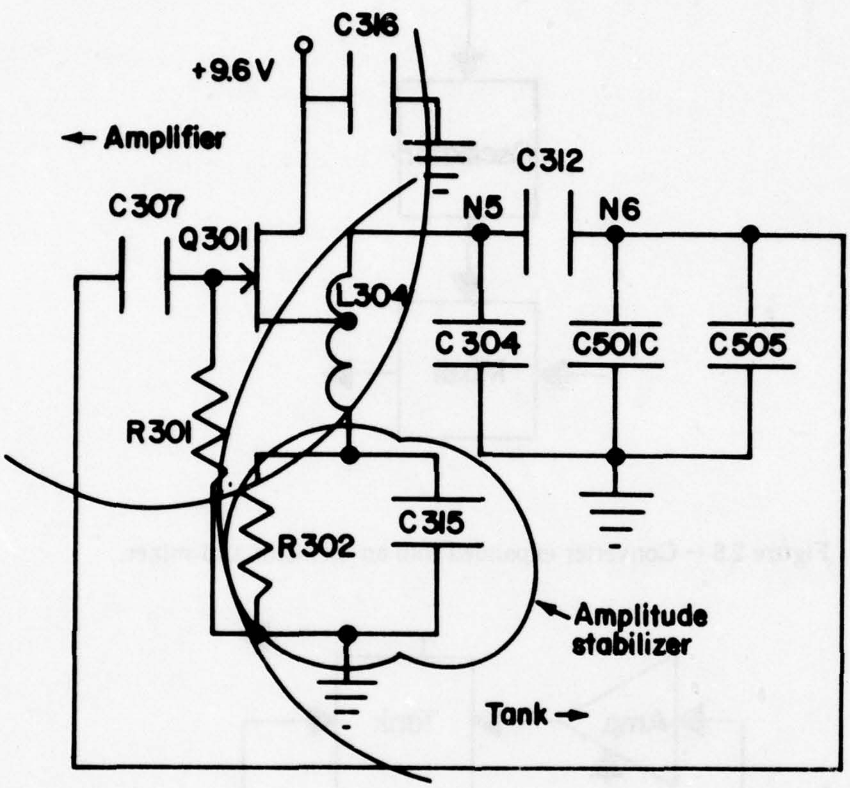

Figure 2.8 -- The oscillator circuit.

and drain biases[1]. What could cause this situation? Remember that WATSON is now looking at the amplifier at the circuit diagram level. Unlike the situation in plans represented by block diagrams, the causal relationships among the parts are not so clear. Knowing nothing else about a particular circuit, WATSON looks for possible failures in the active components first. One such possibility is the opening of the p-n junction at the gate of the JFET. Analysis indicates that the

----

1 If the oscillator were operating, the bias at the gate would be considerably different from the quiescent bias. The mechanism that supports the operating bias scheme is quite difficult to deduce because it depends on the dynamic non-linearity of the JFET. Thus it seems appropriate to represent it explicitly. If it is quiescent, the JFET is biased for class A operation. This enables the oscillator to start up. Figure 2.9 is the relevant subcircuit for this bias scheme. Commentary on the oscillator and its amplifier mentions both bias schemes.

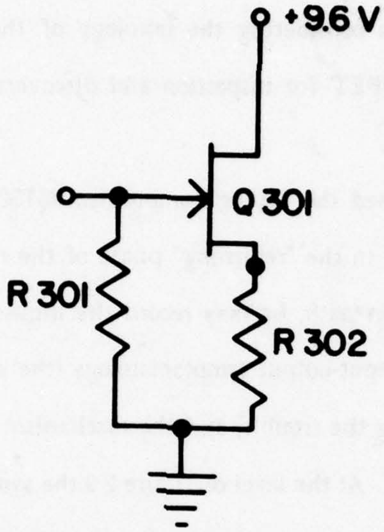

Figure 2.9 -- The DC subcircuit.

observed *drain current* would be larger than the drain current expected (when operating faultlessly), given the observed gate-source voltage. An opened gate implies zero gate current. hence the voltage drop across R302 must be zero. These DC consequences predicted by WATSON (in the context of class A operation) agree with reality.

What about the AC consequences? WATSON reasons that a displacement up or down from the gate bias voltage would have no effect on the output side of the transistor given the rules governing Q301's behavior in the opened gate situation. How does this AC symptomatology at the circuit diagram level *lift* to the block diagram of figure 2.7? WATSON knows that the purpose of the feed-back loop is to stabilize the output amplitude of the oscillator. Would the oscillator remain stable if the amplifier were not doing its job of providing AC gain? The answer is no, for WATSON reasons as follows: Suppose the amplitude at the output of the tank were to decrease. To restore the output to the desired amplitude, the amplifier must kick the tuned circuit that underlies the tank module. But if the JFET is actually failing in the mode proposed, there is no way to make the JFET's gain be non-zero. Eventually all of the oscillatory energy will be dissipated and the oscillator will become quiescent as observed. Notice that the driving force

in WATSON's reasoning comes from considering the teleology of the stability arising from the feedback loop. WATSON pulls the JFET for inspection and discovers that it is indeed opened at the gate[1].

Having successfully isolated the failing component, WATSON abstracts the bug he has discovered. The abstraction occurs in the "returning" phase of the recursive localization process. For each plan along the localization path, he may record the input-output symptomatology that caused him to visit the plan, the input-output symptomatology (the sign) of the part of the plan which was discovered to be causing the trouble, and the mechanism by which the failure of the part causes the observed behavior[2]. At the level of figure 2.9 the symptom is that there is no AC gain, the sign is an unusually low gate to source bias for class A operation. The mechanism is that incremental displacements at the gate cause nothing to happen at the source. At the level of figure 2.7 the symptom is that there is no oscillatory output, and the sign is that there is no incremental gain[3]. The mechanism has two parts. If the circuit were ever oscillating, it would have ceased to do so owing to dissipation in the passive modules which is not offset by the active module. If the circuit never oscillated, it will never start since doing so depends on having ambient thermal noise amplified by the amplifier of the oscillator. For figure 2.6 the symptom is that the expected output mentioned above is not there. The sign is that the oscillator is not oscillating. The mechanism is implicit in the "tuned" nature of the mixer: mixing the input signal

----

[1] I should also point out that we have only seen WATSON's success paths. At each level of hypothesis generation, there may be more than one plausible hypothesis. WATSON attempts to eliminate as many hypotheses as he can in the rationalization process. It is possible, nonetheless, that several components may be equally plausible loci of failure. In this case the real culprit can only be determined by removing the components in question from the circuit and measuring their intrinsic properties.

[2] Certain criteria are applied to decide whether or not a bug is worth remembering. The nature of the criteria will not be understandable until considerably more of WATSON's machinery is displayed. For now, we shall imagine that an abstraction occurs for every level of localization. As it turns out, the description of the bug mechanism will be intimately related to the notion of a part's purpose in a plan.

[3] Note that there may be no reliable way of checking for this sign without breaking the loop that is intrinsic to the oscillator.

with the oscillator signal that is not 455 kHz offset by the broadcast frequency will cause the broadcast signal to be swallowed up, just as ought to be the case with stations to which the radio is not tuned. By composition, the converter fails to pass the input signal as well. For figure 2.5 the symptom is that there is no output, the sign of the converter is that there is no output, and the mechanism is the one that follows directly from the composition of parts: the lack of output from a part of a flow process leads to a lack of output from the overall process.

2.2 A shorted capacitor.

In this scenario, the presenting symptom of the GR-78 is that it has no audible output when tuned to weak stations. Very strong stations are barely audible and heavily distorted. Again WATSON refers to the plan described in figure 2.4. He requests that a signal generator be set up to generate a 5 MHz carrier of moderate strength modulated at 1 kHz while the radio's volume control is set to a normal listening level. The radio is tuned to 5 MHz. He notes the lack of audible output and asks for an increase in signal strength until an audible output and distortion appear. WATSON performs similar experiments at 0.1 MHz intervals over the band. At each "station" the same symptomatology may be induced, and remains present independent of volume control setting.

With the signal generator producing a 5 MHz signal, modulated at 1 kHz, and of sufficient strength to induce distortion, WATSON begins the back-trace. Knowing that the output of the AF section is wrong, he looks at its input and discovers that it too is distorted. Since the input to the RF section is known to be correct, the RF section is conjectured to be the site of the problem. Expanding the RF section into a more detailed plan (again represented by figure 2.5) he quickly discovers that the output of the IF strip is distorted, but its input (from the converter) is not[1].

- - - - - - - - - -

[1] Again this is a difficult RF measurement requiring both a sensitive RF probe and a frequency analyzing instrument. A not so well equipped technician would hypothesize the reasonableness of the IF strip's input from the converter, and continue.

WATSON is not yet ready to blame the trouble on the IF strip. Note that there is an auxiliary input from AGC2 to the IF strip. In fact, the IF strip and AGC2 together form a feedback loop whose teleology is known to WATSON. The loop's purpose is to stabilize the strength of the IF strip's output signal. WATSON's theory of mechanism for this loop is that the output signal strength's rising above the desired value causes the gain of the IF strip to be decreased. Similarly if the output signal strength falls below the desired level, the gain of the IF strip increases. Another feature of the mechanism is that it has an associated time constant. That is, the description of the mechanism explicitly includes the notion that the strength is determined by time-averaging over a known interval. (The formal details of this description are elucidated in section 5.3.)

In the present situation, an examination of the port that brings the signal from AGC2 to the IF strip reveals a bias of 0 V -- clearly incorrect. How does WATSON know this? He deduces it from the teleological commentary associated with the feedback loop together with information about the input signal strength obtained from measurement. Moreover there are caveats associated with the IF strip indicating minimum and maximum values that the bias at the automatic gain control port can take on. The bias at that port is at ground, hence well under the advertized minimum value. What effect does this have on the IF strip? As already mentioned, the IF strip's associated commentary explicitly prohibits its use with the observed biasing. Consequently to understand how the IF strip is affected by the adverse biasing, WATSON must look at its plan structure in greater detail. Forward reasoning on the plan described in figure 2.10 indicates that the output of the IF strip will be at best an amplified and harmonically distorted image of its input -- if the input signal is large.

The plan for AGC2 is illustrated in figure 2.11. Looking at the input of the low-pass filter (LO-PASS), WATSON finds an amplified and rectified version of AGC2's input. The gain of the AVC amplifier is, in fact, not as large as it should be. But back-tracing in the usual fashion, WATSON notices that the input and output of the low pass filter are not identical --

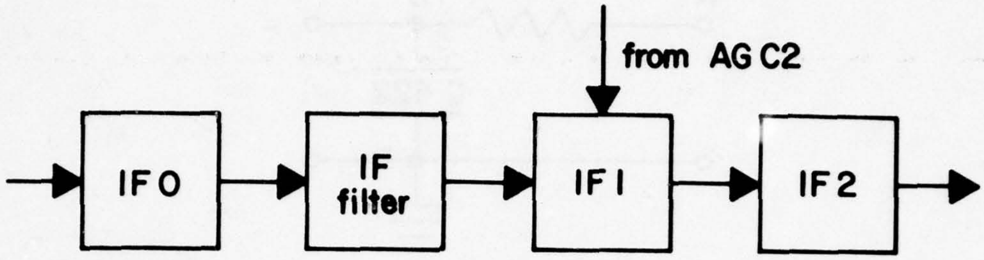

Figure 2.10 -- An expanded IF strip.

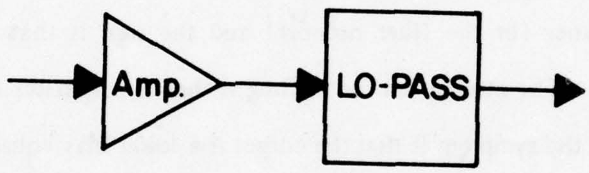

Figure 2.11 -- An expanded AGC2.

modulo a DC offset -- as prescribed. In fact, the input to the filter seems to be wiggling about while the output is quiescent. WATSON becomes suspicious of the filter and expands it into the circuit diagram of figure 2.12. From *a priori* probabilities of failure he posits the shorting of C422. WATSON's model of this network (when operating correctly) is that C422 likes to keep node B from changing. That is, tugging on node A has no immediate effect on node B. Tugging on A for a while will eventually affect B. On the other hand, if C422 were to short, tugging on A could never affect B since the latter node would be at ground potential. This successfully explains the quiescent output of the filter. Is the behavior of the IF strip also explainable by this failure? Yes, since the behavior of the IF strip has already been explained in terms of the bad bias condition. Removing and checking C422, WATSON discovers that it does not meet its intrinsic specifications.

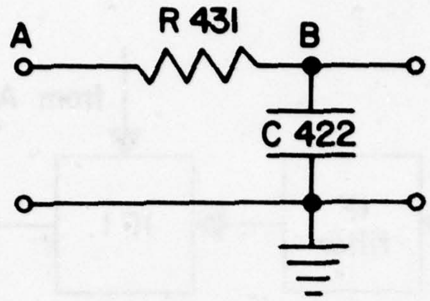

Figure 2.12 -- Circuit diagram of the low-pass filter.

Bug abstraction proceeds essentially as before. For the plan of figure 2.12 the symptom is a low input impedance for the filter network[1] and the sign is that the voltage across the capacitor is nearly zero. The mechanism for the bug is that the capacitor is behaving like a wire. For AGC2 (figure 2.11) the symptom is that the output is a lower bias voltage than expected. The sign is that the AVC amplifier has a lower gain than expected. The mechanism is that the low-pass filter loads the AVC amplifier so as to turn AGC2 into a source of 0 V. The symptom for the the RF section (figure 2.5) is that it has no discernible output, while the sign is a low bias voltage at the port between the IF strip and AGC2[2]. The mechanism is somewhat harder to construct in this case since it cannot be made directly from the parts of the plan for the RF section.

Recall that WATSON had to look inside the IF strip plan (figure 2.10) to decide what the IF strip would do if its control bias were not up to specification. WATSON deduced that especially strong signals would get through (though distorted), while weaker signals would be eliminated.

----

[1] Note that the output port of AGC2 (and consequently the output port of LO-PASS) is a voltage port, hence looking into a high impedance which should be reflected in a high input impedance for the filter network

[2] It is interesting to observe that a broken IF amplifier, due to an open transistor perhaps, would not exhibit the same sign as in the case of the broken AGC2. This is because the input to the IF amplifier from the AVC amplifier would be biased high, attempting to compensate for the broken IF.

The mechanism for the bug in the RF section therefore is that especially strong signals get through. Bias considerations inside the IF strip suggest that the bottoms of modulation signals may be clipped hence introducing distortion into those strong signals that are passed through. Since weak signals are expunged completely by the IF strip, inheritance implies their being expunged by the RF section as well. Abstraction of the bug to the top-level plan (figure 2.4) follows by composition of signal processing parts and inheritance of signal properties at ports shared at various levels in the plan-fragment hierarchy.

2.3 A misaligned front end.

In this scenario the presenting symptom of the GR-78 is that some station is audible in two places on the tuning dial. The plan of figure 2.4 tells WATSON that if the station is audible in one place, it should not be audible in the other. As before the signal generator is set up with a 5 MHz carrier and a 1 kHz modulation. The receiver is tuned to 5 MHz. The 1 kHz modulation is audible. The receiver's tuning control is then "swept" from 5 MHz downward. At a tuning of approximately 4.1 MHz the 1 kHz modulation is heard again without changing the carrier frequency of the signal generator. The plan of figure 2.4 indicates that the 5.0 MHz "station" should be inaudible. Successive refinement of the receiver's sweep setting shows that the output signal strength is maximized if the second tuning is set at 4.09 MHz.

With the experimental set-up described above still running, WATSON begins the back-trace. He finds that the RF section output (AF section input) is a strong 1 kHz signal -- which is wrong. Presumably the problem is not in the AF section since its input is a 1 kHz audio signal. Referring to the plan of figure 2.5, WATSON discovers a 1 kHz modulation component at the detector input and the IF strip input (converter output). The bad converter output is reasonable given that it input contains an RF carrier at 4.09 MHz[1].

----

[1] This conclusion is obtained by expanding the converter plan and noting that the mixer is just as happy to mix down the *image* of the signal as it is to mix down the signal itself.

The *trouble, then, seems* to be in the RF amplifier. The teleological commentary on the RF amplifier indicates to WATSON that it serves two purposes -- to provide sensitivity and to provide selectivity. The first is achieved by amplification. Observation indicates that RF amplifier is successfully fulfilling this obligation. The second purpose includes providing a particular kind of selectivity called *image rejection*. This is achieved by passing a spectrum centered around the broadcast frequency with bandwidth significantly less than the IF frequency[1]. By observation the RF amplifier seems to be passing a significant image signal component. More explicitly it is failing to fulfill one of its design obligations. Hence LS3 suggests looking into the RF amplifier.

WATSON expands the RF amplifier into the plan of figure 2.13. He attempts to carry out a back-trace, but the measurements expert reports that back-tracing cannot produce reliable information since probing the inputs and outputs of the filters will result in changing their pass-band properties. WATSON must therefore pursue another failure hypothesis strategy. He notices that the two filters synergistically accomplish image rejection by successive "distillations" of the 4.09 MHz signal and successive "concentrations" of the 5 MHz signal. What relations must hold in order for this distillation process to work? The answer is that the two filters, which are narrow band-pass elements, must have the same center frequency (modulo stagger or offset, perhaps). (This chain of reasoning is really quite straight-forward, as the detailed account in sections 6.2 and 6.3 will reveal.)

In the present case the center frequency of each should be 5 MHz. Commentary requiring the filters to agree on center frequency suggests that *misalignment* as a plausible cause of the observed symptom. WATSON invokes the alignment expert to look into this problem, advising the expert that the two filters seem to disagree about the frequency of interest and that there are controls that are accessible to adjust the alignment of the filters. Note that the

----

[1] The bandwidth is actually much less than the IF frequency since the filters are known to be moderately high "Q" tuned circuits.

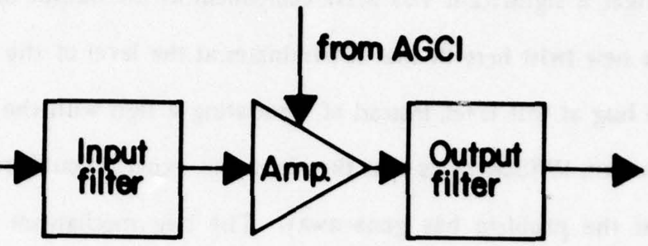

Figure 2.13 -- An expanded RF amplifier.

invocation of this expert was triggered by encountering a problem fitting the following general description: There is a sequence of band-pass parts whose purpose is to reject some frequencies and retain others. Their center frequencies' disagreeing would explain the observed symptoms. Moreover the overlap of those frequencies is adjustable by externally available controls. Note also that the alignment expert is non-local in the sense that it will probably have to appeal to plan knowledge unavailable in the plan within whose scope the expert was invoked. The loop breaking expert is non-local in this sense as well, since handling of loop structures requires simultaneous access to the details of the plans associated with the parts of the loop, and to the overall loop structure.

The alignment expert considers a "flattened" plan of the GR-78. This plan is obtained by manipulating the plan binding hierarchy that underlies the design, causing the innermost plans having band-pass commentary to appear at the top level. The alignment expert performs the usual task of first tweaking up the radio with respect to the high end of the band, then the low end. With respect to each end, the expert first sets the dial calibration by adjusting the oscillator (to yield maximal signal strength) and then proceeds along the signal path from antenna to detector adjusting each of the (adjustable) band-pass elements. This process is iterated until no improvement is achieved.

When the alignment expert finishes, WATSON looks to see if the problem has gone away.

Indeed there is no longer a significant 4.09 MHz component at the output of the RF amplifier. Bug abstraction has a new twist here in that it terminates at the level of the RF amplifier plan. When abstracting the bug at this level, instead of associating a sign with the symptom of strong image signal at the output, WATSON notes that the alignment expert should be invoked, followed by a verification that the problem has gone away. The bug mechanism is recorded as the enhancement of the image frequency with respect to the signal frequency. Bug abstraction at higher levels of the plan hierarchy goes through in much the same way as previously.

2.4 An opened collector-base junction.

The symptomatology of this last scenario is that the radio seems to have no audio output. After some preliminary checking to verify that the radio has no audible output, independent of station selection and volume setting, WATSON calls for the usual set-up. Indeed, with the signal generator producing a 5 MHz carrier modulated at 1 kHz, the radio does not show a 1 kHz output. Referring to the plan of figure 2.4, WATSON determines that the inputs of the AF section meet specifications, whereas its output does not. Hypothesizing that the AF section is the culprit, WATSON pulls out its plan as represented in figure 2.14. The usual measurements indicate that the inputs to the audio amplifier too meet specifications, but the output does not.

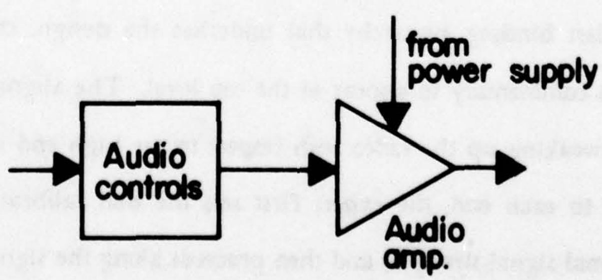

Figure 2.14 -- An expanded AF section.

Expanding the audio amplifier, WATSON finds the plan of figure 2.15. Tracing reveals

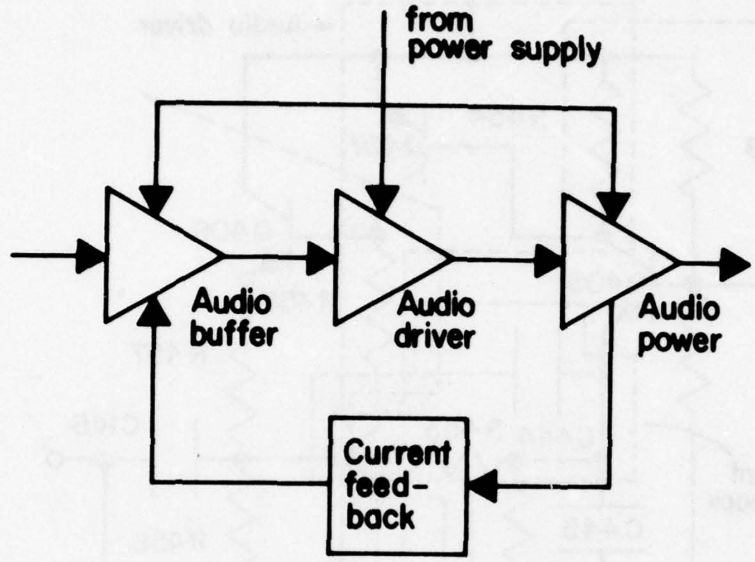

Figure 2.15 -- An expanded audio amplifier.

that the input to the audio power amplifier is quiescent -- which it ought not be[1]. A current probe placed at the input to the audio driver reveals a reasonable input signal[2]. WATSON therefore hypothesizes that the audio driver is the cause of the trouble and pulls out its plan. In order to understand WATSON's reasoning that will lead to localizing the failure, we will have to consider the circuit (as a whole) that underlies the audio amplifier plan of figure 2.15. There are of course plans for each of the parts in the block diagram of figure 2.15. The "dashed" boundaries encompassing the circuit components of figure 2.16 are approximations to those plans.

----

[1] The base biases of its transistors are also wrong in that they are pinned to ground, though the observation is not made at this level of detail. In signal back-tracing DC biases are typically ignored.

[2] The audio driver is current driven by its input signal, hence it is hard to measure that signal with great accuracy. In the present case a current probe will give sufficiently reliable information to determine whether or not the audio driver input is behaving reasonably, i.e. moving the voltage at the base of Q406 toward 9.6 V will cause the current into the base of Q407 to increase markedly. (See figure 2.16.)

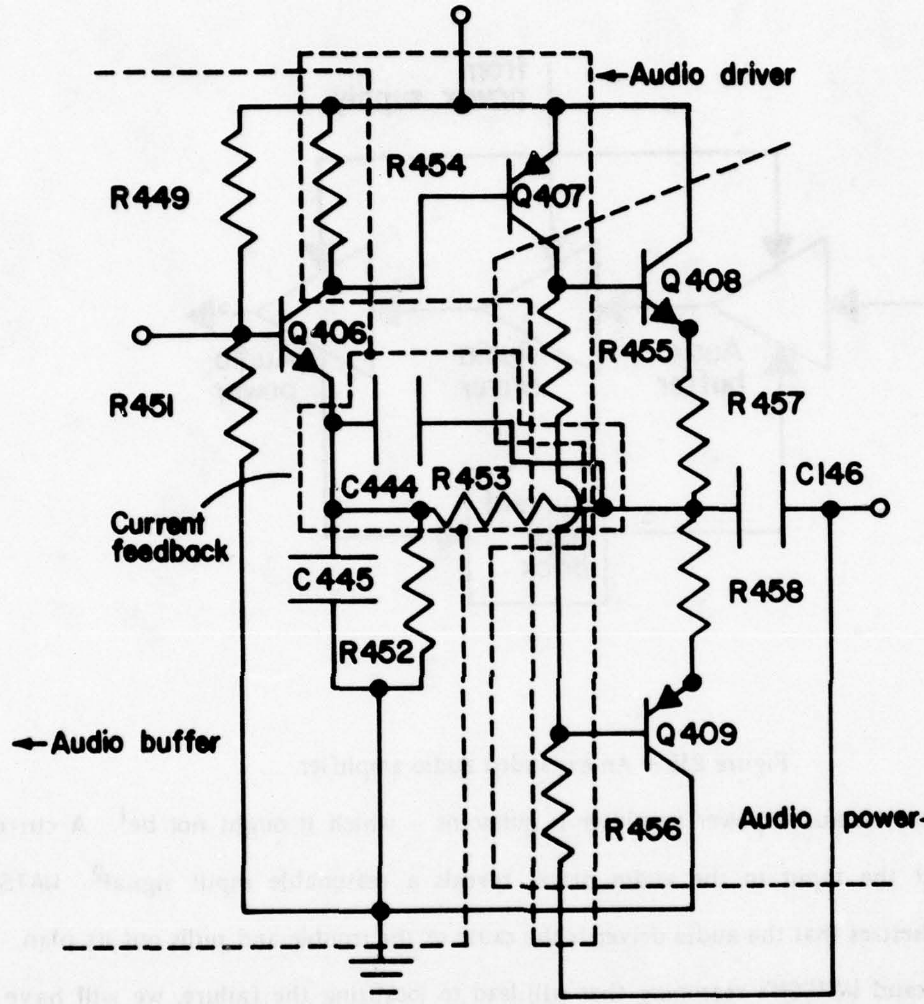

Figure 2.16 -- Circuit diagram underlying the audio amplifier.

WATSON wants to explain how the audio driver might have gone sour. Since he is looking at a circuit plan, he resorts to hypothesizing troubles in active components, because he

knows that their *a priori* probabilities of failing are greater than those for passive components. Consequently he focuses on Q407, a bipolar junction transistor (BJT). The most common way for such a component to fail is by the opening of one or the other of its junctions. From WATSON's point of view they (the opening of one or the other junction) are both equally good hypotheses, so he tries both, hoping to eliminate one or the other during the rationalization process.

This process goes forward in two phases. WATSON first predicts the DC consequences of the hypotheses and checks to see if the predictions correspond to reality. He then does a similar prediction and check for the AC consequences. Under the hypothesis of the open emitter/base junction the DC current in the collector/emitter branch of Q407 should become small. Now WATSON knows that Q407 biases its collector by injecting current into the series resistance composed of R455 and R456. If the collector/emitter branch current falls, the bias voltage at the collector of Q407 (base of Q408) falls as well. In the plan for the audio power amplifier R455 and R456 are used as a voltage divider that set the base bias for Q409. Since the voltage at the top of this divider falls under WATSON's present hypothesis, the voltage at the center falls as well. Thus the base bias of Q409 also falls toward ground.

The transistors Q408 and Q409 comprise a complementary symmetry pair. Under normal circumstances they are just barely turned on. If their base biases fall toward ground, they are completely shut off, though in falling, Q409 will at some point be turned on hard. In particular, neither transistor would have significant current in its collector/emitter branch. Hence the common node of R457 and R458 would fall toward ground too. This results in less current flowing through the feed-back resistor, R453. Therefore less current flows into the emitter node of Q406 from this network. Notice that Q406 is operating in its active region. In particular that means that this silicon BJT will produce as much current in its collector branch as is necessary to keep its emitter within 0.6 V of its base. Hence the current in that branch must increase to compensate for the loss of current formerly coming from the feed-back network.

Increased current in the collector/emitter branch of Q406 means more current coming

out of the collector node into the transistor. Symmetrically this means more current coming into the node from R454 and Q407. But WATSON has hypothesized the opening of the emitter/base junction! So the increased current must come entirely out of R454. This would result in a significant drop in the bias voltage of Q406's collector. Actual measurement indicates that this is not the case. This means that the hypothesis of the open emitter/base junction cannot be right. An alternative hypothesis, the opening of the base/collector junction of Q407, is tried. This leads to identical deductions up to and including the increased current in the collector/emitter branch of Q406. Again WATSON reasons that the increased current must come out of R454 and Q407. But now the emitter/base junction of Q407 is still a forward biased diode. Consequently the base of Q407, a germanium BJT, remains within 0.3 V of the emitter. So the bias at the base of Q407 remains fixed and the increased current actually comes out of Q407. All DC predictions are verified by measurement.

Now WATSON carries out the AC phase of the predictive process. The analytic method is not unlike that used to predict DC consequences. WATSON knows that a current encoded audio signal is presented at the input port of the audio driver. In order to understand the effect of such a signal on the audio driver he imagines the effect of positive and negative increments in current around the base bias current of Q407. He thinks of these increments as rising and falling at rates consonant with the 1 kHz audio signal that is present at the input port of the audio driver. To get the flavor of the analysis let's first look at the case of a normally operating Q407.

An incremental increase in current into the base of Q407 would result in an incremental increase in the voltage at the base of Q408 and an incremental increase in the voltage at the base of Q409. (Remember that R455 and R456 form a voltage divider.) This turns on Q408 a little, resulting in a larger (in magnitude) current flowing in its collector/emitter branch. This in turn means an incrementally larger current flowing to the right through the coupling capacitor C446. This has a positive feed-back effect on the voltage divider. That is, the base of Q408 is pulled

up even more. This is the mechanism by which Q408 is jolted into its active region. Symmetric reasoning on an incremental decrease in current into the base of Q407 leads to deducing an incremental increase in the current flowing toward the left through C446. This reasoning on voltages and currents lifts to the signal level as a voltage encoded signal at the output of the power amplifier that tracks the input to the audio driver. AC analysis must take care to note coupling capacitors like C446 and bypass capacitors like C445. Otherwise the local reasoning processes are indistinguishable from those applied in DC analysis.

Now to get back to the AC analysis of the hypothetical failure: An incremental increase in the current at the base of Q407 results in no change at the collector node. Similarly, an incremental decrease in the base current also results in no change at the collector. In terms of the signal at the output of the audio driver, this AC analysis would indicate no observable signal -- which is precisely the complaint. Pulling Q407 out and examining it shows that its collector/base junction has indeed opened, verifying the hypothesis.

In abstracting the bug for the audio driver, WATSON associates the symptom of a quiescent output in the face of an active input with a sign that includes the various bias changes that were predicted -- and observed -- in the audio driver. The mechanism follows directly from lifting the AC voltage/current behavior to the level of the abstract signals at the input and output of the audio driver amplifier. Abstractions at higher levels of planning are simpler compositions of signal processsing.

3 A Preview of the WATSON Program

Before delving into the details of the machinery enabling WATSON to carry out his task, we should stand back and look at his overall structure, developing a picture both of the major parts of his anatomy, and of their interactions. Figure 3.1 shows a flow diagram for WATSON. In that time-worn tradition, polygons denote essentially procedural structures, and balloons denote bodies of knowledge. Control flows along the solid arrows, while information flows along the dotted arrows. This diagram is somewhat fictitious, first because it is an incomplete representation of WATSON's control structure. Second, the segregations suggested by the closed figures are conceptual and do not necessarily correspond to independent realizations within WATSON. Nevertheless, the flowchart presents many of the essentials of the CONNIVER and LISP functions forming WATSON's control structure, together with the CONNIVER data base in which his knowledge resides.

3.1 Stating the problem.

WATSON accepts complaints about some particular radio design. A complaint is formulated as a pairing of inputs and outputs. The intended interpretation of this pair is that the outputs are incorrect given the inputs. A complaint is formalized as an s-expression (read by WATSON) of the form[1]

```
(COMPLAINT
  design
  list-of-input-signals
```

----

[1] Syntactic variables will be indicated by the use of the lower-case Roman font. Optional structures will be enclosed in "chevron" angle brackets. Alternatives will be indicated by vertical bars.

list-of-ouput-signals
«control-variable-bindings»).

The COMPLAINT may variabilize various properties of the input and output signals, declaring the kinds of values these properties may take on. WATSON first asks himself if the COMPLAINT is reasonable. The determination of reasonableness is made by first matching the inputs mentioned against whatever expections for such input signals might be specified in the design. If the inputs fail to match such specifications, the COMPLAINT is ill-founded, and WATSON so informs the plaintiff of his error. Otherwise WATSON proceeds to determine whether or not the mentioned outputs jibe with the mentioned inputs. This latter determination is made by first inferring the expected outputs of the radio receiver from the reported inputs by the use of various rules. The inferred outputs are then matched against the outputs reported in the COMPLAINT. If the match succeeds, the COMPLAINT is ill-founded. Otherwise, the COMPLAINT is valid, and WATSON retains a record of the mis-match so as to facilitate the construction of a test bench set-up (for debugging purposes).

COMPLAINT validation is somewhat complicated by the fact that a COMPLAINT need not completely specify the operating configuration of the faulty receiver. Thus external controls, for example, may provide a number of degrees of freedom in operating specifications which must be pinned down when debugging. In fixing these degrees of freedom (which could have been specified completely in the COMPLAINT) WATSON may conclude that the COMPLAINT is ill-founded because of unreasonable configuring of the receiver's controls. For example, a COMPLAINT suggesting that a receiver has no audible output, but not specifying the volume control setting, may be determined to be ill-founded on the grounds that when the volume control is set at a reasonable level, the receiver does have an audible output. The operating configuration for the debugging procedure is actually determined in the same operations as validation of the COMPLAINT with respect to the settings of external controls. The fixing of the operating configuration is a consequence of the process of establishing the validity of the COMPLAINT under whatever constraints the COMPLAINT imposes.

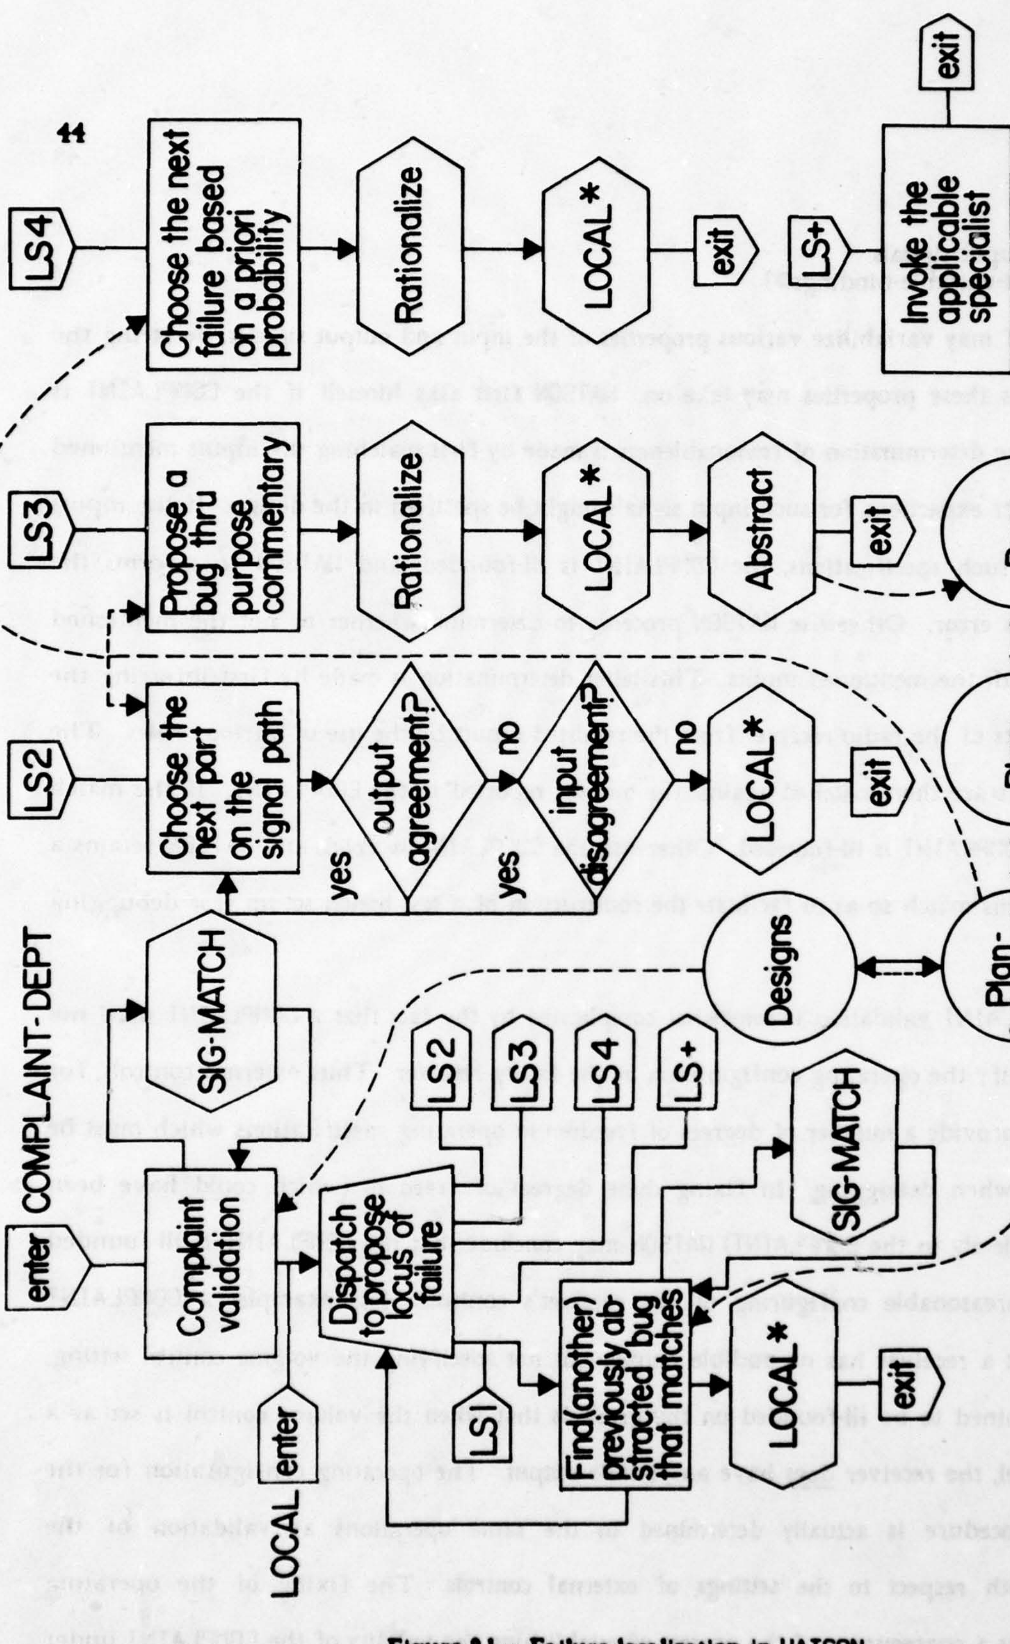

Figure 3.1 — Failure localization in WATSON.

45 Localizing the failure.

When a COMPLAINT is lodged, WATSON is entered (see figure 3.1) at the label, COMPLAINT-DEPT, and is validated as per the discussion above. Then the real work begins with a call to LOCAL. As mentioned previously, WATSON has available a number of localization strategies, but as with all reasonable problem solvers, he always attempts to use the strategy which asks, "Is the answer already known?" Hence, he first dispatches to LS1 to see if a bug has been previously abstracted that covers the present situation. The principal determiner of the applicability of an abstracted bug is the general signal matching processor, SIG-MATCH, but as we shall eventually see, other recognition criteria may be in order as well. If a bug fits, WATSON is left with a sub-plan-fragment (of the plan-fragment to which LOCAL is currently being applied) which contains the locus of failure.

The dispatch mechanism is not a simple COND, of course, for the strategies dispatched to may be exited at various stages of completion. LS1 may run out of applicable bugs, for example. Or the successful localization at one level may not yield a successful localization at the next level. Hence the hypothesis implied by a localization (that a particular sub-plan-fragment is at fault) would have to be suspended to admit trying some other possibility. The dispatch is further complicated by the fact, for example, that the back-tracing strategy, LS2, may make immediate reports to the current activation of LOCAL about interesting phenomena it (LS2) encountered. Such reports may cause the suspension of the back-trace and the activation of some other strategy.

Note that LS1 relies on the body of knowledge in the 'plan-fragments' balloon. That is not to say, however, that that is the only source of information. The double-headed solid/dotted arrows indicate that information from one body of knowledge may be fetched (or added) through another such body. The mechanisms for doing these indirect acquisitions of information are typically embedded as CONNIVER methods in the data base.

If no applicable bugs are known, WATSON looks around for another localization strategy, one of which is LS2, the back-tracing strategy. This strategy is applicable to cascade plans, and is

an iteration on parts found along the signal path of such a plan. The basic iterative step is to examine whether or not the observed output of a part agrees with the predicted output (given the parts observed inputs). If a disagreement is detected, the strategy next looks to see if the expected and observed inputs agree. If they do, the part "sandwiched" by those inputs and outputs is presumed to be at fault. The polygon labelled 'LOCAL' suggests a recursive call (denoted by the star) to LOCAL. LOCAL is applied to the plan-fragment corresponding to the "found" part. Once again, I should like to emphasize that the flowchart omits a great deal about the flow of control in WATSON. Several kinds of failure are possible on the LS2 branch of the localization process. Such failures can cause a number of different changes in the flow of control, ranging from the temporary suspension of a hypothesis (of failure of a part along the trace path), to complete abandonment of the LS2 branch. LS2 draws principally on knowledge from the 'plan-fragments' balloon.

The LS3 strategy is typically entered by virtue of LOCAL's having received a report from some other strategic branch indicating that some well-defined difficulty has been encountered. "Well-defined-ness" is determined by the failure to meet some criterion of the design. This suggests checking to see where the responsibility lies for meeting such a criterion. If a sub-plan-fragment with the appropriate purpose commentary can be located, an investigation ensues to see if the failure to meet the responsiblity (by the identified part) could lead to the observed misbehavior in the local plan-fragment. If this rationalization is successful, a process of abstracting a description of the bug and its underlying cause ensues, and a record of the abstraction is made in the data base under the heading 'bug' as a note on the plan of which the local plan-fragment is an instance.

LS4, the strategy based on *a priori* probabilities of failure, is only applied when WATSON has come within immediate reach of resolving the underlying cause of the initiating COMPLAINT. This is the case when the problem has been localized to a plan of the circuit or coupling type. At that point sub-plan-fragments will typically be instantiations of plans of the component type.

hence will be annotated with descriptions of their failure modes, including likelihood of occurrence. The failures are sorted in the order implicit in the annotation, are selected from in that order, and are rationalized with respect to the observed AC and DC behavior of the circuit (coupling) plan. A successful rationalization will lead to a recursive invocation of LOCAL. This invocation will terminate the localization process if it is applied to a plan-fragment corresponding to a component, resulting in the pulling of that component for inspection of its intrinsic properties.

Finally there is the LS+ branch of possible localization strategies. This is the catch-all strategy which invokes various specialist experts based on the encountering of particular kinds of structural impediments to the other localization strategies. There is presently only one such expert contemplated for implementation -- the loop breaking expert. Others, however, are easily imaginable. For example, it might be appropriate to have an expert to deal with 60 Hz hum, a common enough bug in radios. When this phenomenon is encountered, there is usually no particular source identifiable for it. In any event experts appropriate to the local structure of the localization problem will in turn call upon other control and knowledge structures in WATSON, including LOCAL.

3.3 Describing radios to WATSON.

Besides the failure localization subsystem, WATSON's other major subsystem is that concerned with plan description and design assembly -- which we shall call the 'assembly subsystem'. The interesting aspects of the control structure of this subsystem would not be particularly elucidated by a flowchart, hence I have chosen not to present one. The major functions of the assembly subsystem include the definition of plans for various functional units in a radio receiver, the compilation of instances of plans as plan-fragments, and the association of plan-fragments in the hierarchy of a radio design. The first function is carried out in a very

straight-forward fashion. The second function -- compilation -- introduces complexity due to the fact that much of the essence of a plan-fragment is procedural in nature. These procedures are stated in a stylized form as rules associated with a plan. The application of a rule may be triggered by various combinations of events (items ADDed to the CONNIVER data base). The particularization of a rule to a plan-fragment turns out to be a rather complex process involving the creation of many CONNIVER functional closures.

The integration of plan-fragments into a design is complicated by the fact that design is typically carried out in a "top-down" fashion that delays the filling in of various "slots" specified in a plan. The delay is due to the fact that the slots can be filled explicitly by subsequent introduction of design sub-struture in the assembly process. Similarly the filling in process may happen implicitly at run-time. By 'run-time' I mean the point at which the design is used to do reasoning about the expected behavior of the radio receiver. As it turns out, all such delayed references are resolved using CONNIVER IF-ADDED and IF-NEEDED methods. Needless to say, the result of entering the assembly subsystem is the adding of facts to the data base under the balloon headings 'plans', 'plan-fragments', and 'designs'.

3.4 Programming constructs.

In the succeeding chapters I shall frequently exhibit fragments of CONNIVER or LISP code. I shall take the liberty of assuming familiarity with the primitives of those languages, as they are well documented elsewhere [McCarthy, 1965; McDermott, 1974b; Moon 1974]. Functions that are peculiar to WATSON will be mentioned as such and will have their semantics explained (if not explicitly, then by context). Most such code fragments will make reference to the CONNIVER data base. Whenever possible, I have tried to avoid the cumbersome CONNIVER pattern matching syntax by using the lower-case Gothic font for pattern-variable names. The nature of the intended match should be clear from the textual context.

There are three syntactic constructs introduced and used extensively: FINDs, FORs, and *path names*. A FIND has the general form

(FIND quantifier variable-list pattern),

where 'quantifier' is a quantifying expression over the list of variables, 'variable-list'. The meaning of the quantifier will always be self-evident. The basic idea is that the mentioned pattern (possibly containing match variables bound in the scope of the FIND, possibly containing match variables free with respect to the FIND), should be matched against in the data base under the constraints implied by the quantifier. For each successful retrieval of a matching data base item, the outstanding bindings of the bound variables are appended to a list. The list is returned as the result of the FIND after all possible matches are made. It will be my habit to mention variables bound in the scope of the FIND as if their bindings were still available on exiting the FIND. I shall do so however only when context makes my intentions unambiguous.

A FOR, which has the form

(FOR quantifier variable-list pattern prog-body),

is much like a FIND. Instead of returning the bindings resulting from successful matches with the data base, for each successful match, the sequence of code 'prog-body' is evaluated. Of course, the code may mention the variables bound by the FOR, which will have values appropriate to the most recent match. I should mention that patterns in both FINDs and FORs will usually be QUOTEd (since that argument position is passed by value). Typically, however, instead of the usual LISP quotation, I shall use the CONNIVER skeleton construct '!"'.

Finally, though all the objects which the assertions in the data base are about generally have canonical names, such names are not exactly rife with semantic content. For clarity I shall generally make use of the path-name construct

~(NAIL INDEX-FINGER RIGHT-HAND G00069),

which denotes the obvious part of a distinguished individual, G00069. This particular path-

name evaluates to the intended canonical object at read-time. There are also other path-name macro characters -- ',~' and ',,~' -- which cause path-name evaluation at EVAL time and pattern match time respectively.

4 The Case of Q301

Having seen what WATSON can do, we are now in a position to look into how he does it. The rest of this exposition will be devoted to explaining the "how." We will accept WATSON's initial knowledge as *fait accompli*, ignoring for the time being the intricacies of the readers and evaluators that make it possible to communicate to WATSON the large data base that encompasses plans, designs, component descriptions, circuit diagrams, etc. We will, however, be much concerned with the internal representations that these data assume, and the manipulations that are applied to them.

Another aspect of the communication problem that we will ignore here is the digestion of the sensory data provided by measuring instruments. These data may be as trivial as a reading from a voltmeter, or as complex as a signal trace on an oscilloscope face. Neither of these is the real object of the measurement activity. The real object is the association of values with the obs of an abstract signal description. Ideally WATSON would be able to take these basic data and translate them. He does not in fact do this. WATSON's measurements specialist, MAXWELL, to which I shall make frequent reference, presents the human assistant with questions phrased in terms of the obs needing value assignments. Hence the burden of translation from sensory data to symbolic description is carried by the human assistant.

4.1 Lodging the complaint.

Recall that in the first scenario (section 2.1) that the presenting symptom was that the GR-78 showed *no output*. The formal statement of this complaint is

(COMPLAINT GR-78

```
(INPUTS
  (~(OBS-SIGNAL PORT-1)
    (CARRIER-COMPONENTS
      ((MODULATION AM)
       (CARRIER-FREQ x)
       (MODULATION-COMPONENTS
         ((MODULATION-FREQ y ϵ {*}))))))))
(OUTPUTS
  (~(OBS-SIGNAL PORT-2)
    (CARRIER-COMPONENTS
      ((MODULATION AM)
       (CARRIER-FREQ 0.)
       (MODULATION-COMPONENTS
         ((MODULATION-FREQ z ϵ {*})
          (AMPLITUDE 0.))))))))
(CONTROL-BINDINGS
  (TUNING x ϵ {*}))).
```

A COMPLAINT is lodged about a particular design, the GR-78. Unhappily, the syntax of COMPLAINTs is quite complicated[1], but the gist of it is that we are displeased with the observed output at the port named PORT-2[2] given the observed input at port PORT-1, and given that the GR-78's external controls are set up in a certain way. WATSON's first order of business is to validate the complaint experimentally and to use the experimental results to design a test set-up for tracking down the bug. To do an experiment, settings of the signal generator controls and the radio must be chosen. The "declaration"

...CARRIER-FREQ x...

tells WATSON to name the carrier frequency parameter of the generated signal 'x', but offers no advice as to how it is to be assigned. The CONTROL-BINDINGS clause specifies which of the radio's control variables -- TUNING (i.e. radio's station selection control) -- is to be affiliated with the carrier of the input signal, x, and specifies that it can be assigned any legal value. WATSON

----

[1] WATSON thinks of an audio signal as a modulated DC signal.

[2] Many of the abstract objects of WATSON's knowledge base -- ports, nodes, plan-fragments, etc. -- are realized by atomic objects called obs with GENSYMed canonical names of the form FROB-105. Such obs can generally be gotten at in either of two ways, by the structural route -- as exemplified by the kinds of data base items we shall be looking at presently -- or by path name. The "gritch" character, '~', is a reader macro character indicating that the subsequent list-structure denotes a particular ob.

finds the range of possible values for that control variable. He can do that because the data base contains the items[1]

```
(PLAN RECEIVER CASCADE)
(PF PF-3 RECEIVER GR-78)
(PF-PART NIL NIL PF-3)
(PF-CTL PF-3 TUNING CTL-4)
(VALUE-RANGE CTL-4
   (CONTINUOUS-GENERATOR (FROM 3000000.) (TO 7500000.)))
(OPTIMAL-VALUE CTL-4 5000000.).
```

Now knowing that the value range is continuous between 3 MHz and 7.5 MHz, WATSON decides to do experiments for 3 MHz, 5 MHz, and 7.5 MHz settings of the radio's station selector.

The input signal described in the COMPLAINT also has the following declaration

...$y \in \{*\}$...

which says that the value of MODULATION-FREQ should be set to whatever is convenient for making the generated signal match the expected input signal for the GR-78. Not unlike the optimal value (the frequency at which the plan works best) seen above, the expected input signal has an optimal value which is identified with the audio mid-band frequency of the radio's AF section, 1 kHz. There is also commentary indicating that the modulation frequency can vary continuously from 100 Hz to 5 kHz. WATSON therefore chooses experimental settings for the modulation frequency of 100 Hz, 1 kHz, and 5 kHz respectively. z, like y can be set to anything that is convenient. Since it is an output parameter, however, 'anything' means that it can potentially take on all possible values. That is, every possible modulation component on the output has negligible amplitude[2]. This exhausts the degrees of freedom mentioned in the input signal.

---

[1] Some of these items are not actually PRESENT in the data base in the formal CONNIVER sense. They may appear by virtue of various deductive methods. It will suit our purposes for the time-being to imagine that they are PRESENT.

[2] This follows from WATSON's essentially linear model of the signal processing carried out by radios.

The next order of business is to use those degrees of freedom implicit in the controls of the radio. Evaluating

```
(FIND ALL (x) '(CONTROL (T-L-PLAN-FOR GR-78) x))
```

yields a list

```
((PF-CTL PF-3 TUNING CTL-4)
 (PF-CTL PF-3 VOLUME CTL-5))
```

of possible bindings for the match variable, x. The first item has already been taken care of by virtue of having generated the three possible station settings. The second item is completely free. A query for the value range of CTL-5 reveals to WATSON that VOLUME is continuously variable over some range, a pair of numbers corresponding to the "full-scale off" and "full-scale on" positions of the volume control knob. He chooses the end-points and middle value for the purposes of the validation experiments.

The result of all this is that WATSON has generated a number of possible configurations for the signal generator and radio by the obvious combinatorics. Actually there are two degrees of freedom, that have not been considered in configuring the experiments. It so happens that the generator's amplitude and percent modulation are selectable. It is WATSON's heuristic inclination to leave these at *a priori* fixed values unless these signal properties (amplitude and percent modulation) are mentioned specifically in the COMPLAINT. Also, the GR-78 design has an input port not mentioned in the COMPLAINT, the power port. WATSON knows about three kinds of ports, those with control signals, those with information signals, and those with power (or bias) signals. WATSON proceeds to validate the complaint in the scope of a loop whose rounds are determined by the configuration combinations developed above. (There are of course, twenty-seven such inner-most rounds.) Forward reasoning is required to deduce what the output should look like given the inputs established by a particular test configuration. The forward reasoning is actually carried out in a CONNIVER context[1] that is private to the current inner-most round. Figure 4.1

---

[1] Since I shall have the occasion to use the word 'context' in a number of ways, 'context' in the lower-case Gothic font is reserved for the CONNIVER usage of the word.

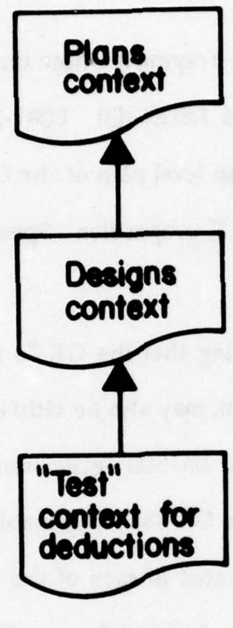

Figure 4.1 -- The context structure at the time of validation.

illustrates WATSON's use of the context machinery. At the top-level is a context containing the abstract plans along with various facts. Below that is a context containing the GR-78 design and whatever facts are local to that design. Finally there is the context in which the deduction will happen.

How is this forward reasoning done? Answering this requires some explanation of the theory of representation embodied in WATSON. We should first realize that most of the items in his data base are propositions over a class of objects called obs. obs are carriers[1] for the attributes represented by the propositions, the latter being realized as CONNIVER items. Earlier in this section we saw the following propositions

(PLAN RECEIVER CASCADE)

---

1 It may also be helpful to think of them as analogous to the formal objects that Sussman [1975] uses in subroutine geralization, or to the anonymous identifiers used by Hewitt [1971] to carry out procedural abstraction, or to Skolem functions [Chang, 1973] in their standard role in the elimination of existential quantifiers.

```
(PF PORT-2 RECEIVER GR-78)
(PF-PART NIL NIL PORT-2)
```

which essentially said that there is a plan- fragment whose unique ob name is PORT-2 and it is a token (at the top-level) of a plan called RECEIVER. PORT-2 is the ob of interest; it is the instatiation of the RECEIVER plan as the top-level plan of the GR-78. One particularly interesting proposition satisfiable by an ob is the VALUE proposition. Specifically,

```
(VALUE PF-3 3500000.)
```

corresponds to the TUNING control indicating that the GR-78 is tuned to 3.5 MHz. The VALUE-RANGE proposition, which we have also seen, may also be satisfied by an ob.

In the validation situation, what WATSON really wants to know is whether or not the signal reported[1] at the speaker port of the GR-78 is reasonable, given the signal reported at the antenna port. Now the input signals generated in each of the validation test cases that he creates are instances of the reported input signal. Let SIG-6 and SIG-7 be the ob names denoting the generated signal and expected signal structures, respectively. Evaluating

```
(SIG-MATCH 'SIG-7 'SIG-6)
```

will tell WATSON whether or not the signals match[2]. If they do, the input is declared to be reasonable, given the design. A successful match leaves the values of the various obs bound, or at least "ranged," in the context of the current experiment.

WATSON matches the reported output signal against the expected output signal However, he will find that many of the obs in the output signal structure, SIG-8, do not have

----

1 WATSON distinguishes signals reported in a COMPLAINT from those that he observes.

2 We need not worry yet about the precise description of the signal structures denoted by SIG-7 and SIG-6, nor about the precise nature of the matching procedure SIG-MATCH. For the time being it should be adequate to think of this latter procedure wandering (in parallel) over a pair analogous structures of a particular kind. The matcher embodies a theory of compatibility for various pairings of ob values and ranges. The procedure reports whether or not visited pairs of obs satisfy this theory. Whenever a matching pair is found, a VALUE-RANGE for that pair is created in the context of the match. A mis-match occurs whenever a VALUE-RANGE is forced to be empty. Hence the matching operation is a test for non-empty intersection.

value properties. What they do have is commentary indicating how such values might be deduced from values of obs of the input signal description, SIG-7. Briefly the value is deduced through the interaction of a number of CONNIVER methods. The first kind of method triggered is an IF-NEEDED -- called a determiner -- whose pattern of invocation matches

```
(VALUE REG-9 value),
```

where REG-9 is an ob whose value is needed for the match of SIG-8 with the observed output signal. The determiner in turn fetches and asserts a number of items. These assertions have the effect of triggering IF-ADDEDs that are the constituents of rules. These latter methods ADD items to the data base that look like

```
(RESULT result-name
   (rule . list-of-results)
   (VALUE REG-9 value)).
```

This says that a kind of fact called a RESULT, whose name is 'result-name' has been derived by the named 'rule' from other facts on 'list-of-results'. The fact in question concerns the value of REG-9. Rules are built up from a complex of IF-ADDED methods whose presence in a context may depend on what RESULTs have already been ADDed to the data base. Thus in evaluating REG-9, the values of its antecedents are asserted (as RESULTs) which will then trigger a rule producing a RESULT giving the value of REG-9. This value finding is recursive in that values for the antecedent obs may have to be gotten by the obvious reinvocation of the same machinery[1]. The io-contour mentioned in section 1.2 is the closure of an ob under RESULTs that determined its properties together with RESULTs it helps determine.

Now WATSON wants the output match to fail. (Remember the COMPLAINT claims a discrepancy between the observed output and the expected output.) The match does fail. Why? Suppose that obs REG-10 and REG-11 correspond respectively to the modulation frequency and

- - - - - - - - - -

[1] In the present case the only way a ob of the input signal description can fail to have a value is by virtue of not having <u>told</u> it to WATSON. Consequently such missing values are found by appealing to the measurements expert.

modulation amplitude of the expected output signal. Let REG-12 and REG-13 be the analogous obs in the observed output signal. Finally let REG-14 be the modulation frequency of the antenna input signal. Via the determiner query above, REG-10, REG-12 and REG-14 are discovered to be equal, as they should be. REG-13 unfortunately has a value of 0, belying the non-zero predicted value associated with REG-11. No match! In fact, any of the twenty-seven experiments in which the volume control is not specified to be at its low end leads to this mis-match. WATSON's final step in each round of the validation loop is to make sure that the measured (observed) signal matches the reported signal. This turns out to be the case, meaning that the COMPLAINT is valid.

Now to establish a test set-up for the localization of the failure, WATSON recollects how the preceding experiments were generated. For all of the continuously ranging control variables (corresponding to certain degrees of freedom in the experiments) he selected high, middle, and low values. High and low values can be gleaned from the salient VALUE-RANGE propositions. The middle value is typically an optimal value (for the performance of the radio), also gotten from explicit commentary, or a mean of the high and the low values. The test set-up will typically be derived from these middling values. In particular, for the rest of the first scenario, the signal generator and receiver are both tuned to 5 MHz with a 1 kHz modulation being imposed on the signal generated. The volume control is at its medium value, and the amplitude and percent modulation of the input signal remain at their "typical" values.

4.2 Signal tracing.

Having satisfied himself that there really is a problem, and having devised a test case which reliably demontrates the problem, WATSON invokes his localization specialist, LOCAL. As we shall shortly see, this a recursive procedure, whose recursive structure parallels the hierarchical structure of plan-fragments embodied in the design of the radio receiver. The first question asked in the localizer is whether there are any known explanations for the observed symptomatology. Evaluating

```
(FIND ALL (bugname sign mp)
  '(BUG bugname RECEIVER
     (SYMPTOMS PF-3
        (INPUTS (PORT-1 SIG-6))
        (OUTPUTS (PORT-2 SIG-8)))
     (SIGNS . sign)
     (MISSING-PURPOSES . mp)))
```

will return all previously abstracted bug descriptions whose symptomatologies match the present case. (Remember that PF-3 is the plan-fragment instantiating the top-level plan for the GR-78.) In the scenarios we have seen, WATSON is presumed to have done no previous debugging, nor is he initialized having abstract bug knowledge associated with plans. Consequently this query in the localizer is unenlightening for WATSON.

Having no *a priori* knowledge about the particular problem facing him, he begins to trace the signal. Understanding this process requires understanding more of the details of of the representation of plans and plan-fragments. We have already seen the items

```
(PLAN RECEIVER CASCADE)
(PF PF-3 RECEIVER GR-78)
(PF-PART NIL NIL PF-3)
```

that tell WATSON that PF-3 is an instance of the RECEIVER plan, at the top-level of the GR-78 design. WATSON also knows that this plan is of type CASCADE, hence it is a reasonable candidate for signal tracing. To do the tracing, however, he needs to know the port/part interconnection description provided by the RECEIVER plan. Items of the form

```
(PORT-SANDWICH RECEIVER part-name1 port-name part-name2|NIL)
(PF-PORT PF-3 port-name can-name)
(PF-PART PF-3 part-name plan-fragment)
```

are sufficient to give this information and to map abstract parts and ports of the plan onto realizations in the design. WATSON already has his hands on the canonical name of the speaker port of the GR-78 -- PORT-2. Three queries of the data base will yield both the parts of the RECEIVER plan that are looking into the speaker port, as well as the names of their plan-fragment instantiations in PF-3.

It is helpful to think of PF-3 as being "active" at the current level of WATSON's recursive localization process. We have already seen that the ports of PF-3 have structures hanging on them that represent the signals at those ports. Such structures are built up from data base items of the form

```
(PF-SIGNAL PF-3 PORT-1 SIG-15)
(MODULATION SIG-15 REG-16)
(VALUE REG-16 AM)
(CARRIER-COMPONENT SIG-15 REG-17)
```

which tell us some of the features of a signal (under a frequency domain interpretation) at the GR-78's antenna port, PORT-1[1]. What I should like to emphasize is the fact that such signal descriptions are local to some plan-fragment -- PF-3 in this case. Why aren't signals globally known as are the ports themselves? Observe that the plan-fragment PF-3 is an instance of the RECEIVER plan, and as such must have sub-structures corresponding to the plan's parts. These parts are associated with other plan-fragments that are instances of yet more plans. Consider then that the hierarchy of plan-fragments, that this organization entails, represents various levels of detail in the design. As WATSON delves deeper into the hierarchy by successively activating various plans, more features of a signal at a given port should become relevant. If plan-fragment A encompasses plan-fragment B (as a part) in such a way that they share a port, there are details of the signal at that port which are appropriate for B to know about, but not A. This problem of local visibility is solved by associating two copies of the signal at the port, one for A and one for B. As we have already seen, part of the strategy for evaluating the obs that comprise the signal is to look for a value on the local copy of the signal, or by determination from its antecedents. If none is forthcoming, the ob evaluator looks up the hierarchy of plan-fragments to see if the value can be found on an analogous ob associated with the signal at the same port. One more

----

[1] Note that the value of the modulation ob is available directly. Again let me emphasize that we will generally assume immediate availability for the sake of simplicity, though such values may actually be found by a determiner or other deductive methods, such as inheritance through the plan-fragment hierarchy.

complication in the representation of signals is the fact that the signal actually appearing at a port may bear no resemblance to the signal that is expected there. The signals that I have been referring to are the expected variety. In parallel with each occurrence of a structure representing the expected signal, there is also a structure representing the real, or observed signal, as is indicated by

(PF-OBS-SIGNAL PF-3 PORT-1 SIG-18).

Now we know enough about the representational machinery to describe the basic signal-tracing localization strategy in some detail. We presume that the strategy is at some arbitrary level in its recursive invocation and have focused on the output port P in plan-fragment PF as having an offensive signal. PF is declared to be the active plan-fragment by binding the CONNIVER variable CULPRIT to PF. Another aspect of activation is the pushing of the old context and the creation of a new one. This corresponds to the fact that WATSON is <u>hypothesizing</u> PF as the source of of the difficulty, a hypothesis that WATSON may wish to back out of at a later time. Such backing out is made very convenient by contexts, since they admit quick dismissal of deductions predicated on the hypothesis that generated the context.

Suppose PF is an instance of the plan illustrated in figure 4.2. WATSON walks along the cascade, internally represented by a collection of PORT-SANDWICH propositions. Beginning at the output port of the rightmost part, WATSON does a SIG-MATCH of the expected signal with the observed signal. Matching against the expected signal may, of course, necessitate propagating (by determiners and rules) the leftmost input signal through the cascade formed by the various plan parts. This process goes as previously described, with the required obs being filled in with their values. The general strategy is to move leftward along the signal path until a part is reached at which the SIG-MATCH succeeds on the output side but fails on the input side. Note that PF must have taken the blame in just this sort of process. Since some part internal to PF shares its good input port, the leftward movement must terminate. This strategy also works in the more complicated case of figure 4.3. The occurrence of a part join requires the tracing process to split

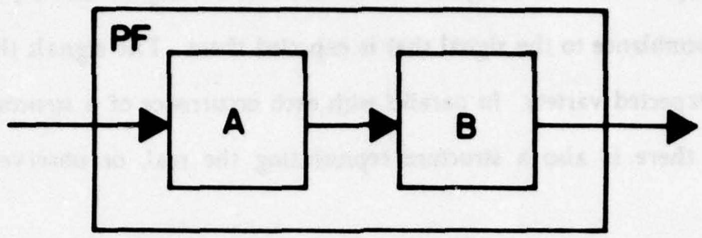

Figure 4.2 -- A simple cascade.

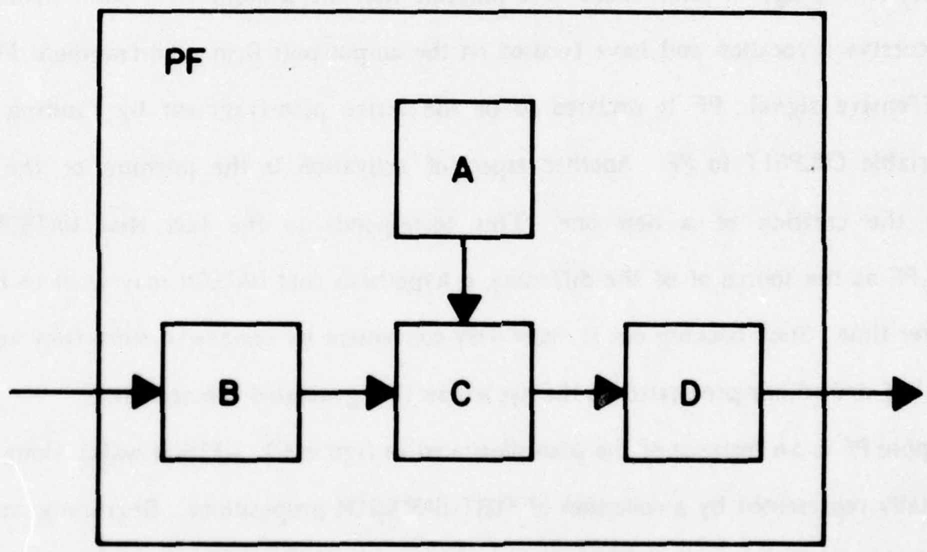

Figure 4.3 -- A cascade with a simple part join.

into two subprocesses. It should now be realized that the set of SIG-MATCHes performed is somewhat more complicated than I actually stated. For each input port of a part, the observed signal is SIG-MATCHed against the expected signal. All such matches must be successful in order to terminate the trace propagation. The trace only propagates at ports showing mis-matches.

Figure 4.4 shows the most general case of tracing topology. The existence of loops in the signal necessitates that the tracing process mark the parts as they are visited. But since there

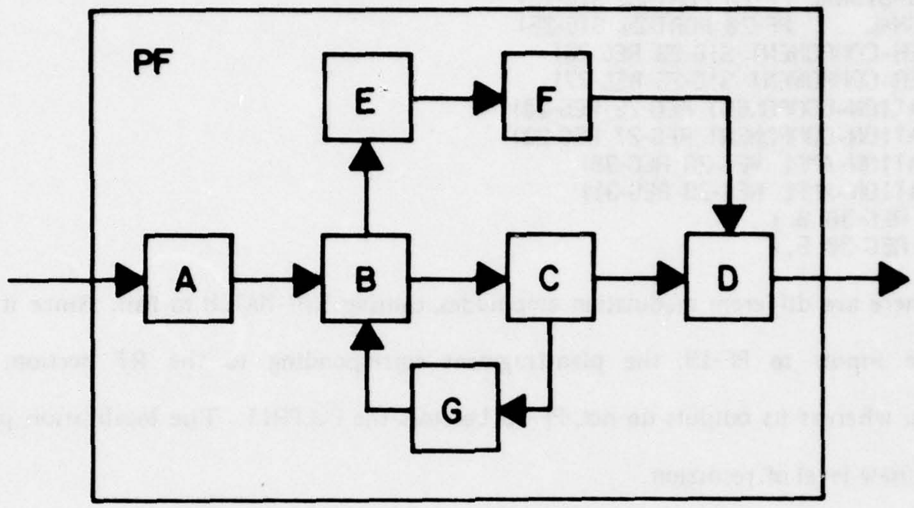

Figure 4.4 -- A fork and join combined to give a loop.

are potentially many instances of the tracing process making visits (consider the fork in the figure) a marking scheme must indicate who made the visit. Things are further complicated by the desire to merge tracing processes whenever possible. Note that only one tracing process need emerge through the input port of B. An ancestral process naming scheme, combined with a "visiting card," indicating the port used by the uniquely identified process in making its visit, solves all of these problems.

A straightforward application of the signal-tracing localization strategy, LS2, underlies the first scenario. The initial recursive application is to PF-3, an instance of the RECEIVER plan. Items in the data base:

```
(PF-PART PF-3 RF-SECT PF-19)
(PF-PART PF-3 AF-SECT PF-20)
```

indicate that this plan has two cascaded parts. Tracing the cascade reveals that the input signal to PF-21, the realization of the AF section in PF-3, does not meet specifications. A partial comparison -- component by component -- of the observed and expected signals

```
(PF-OBS-SIGNAL PF-20 PORT-22 SIG-23)
(PF-SIGNAL     PF-20 PORT-24 SIG-25)
(CARRIER-COMPONENT SIG-23 REG-26)
(CARRIER-COMPONENT SIG-25 REG-27)
(MODULATION-COMPONENT REG-26 REG-28)
(MODULATION-COMPONENT REG-27 REG-29)
(MODULATION-AMPL REG-28 REG-30)
(MODULATION-AMPL REG-29 REG-31)
(VALUE REG-30 0.)
(VALUE REG-30 5.)
```

shows that there are different modulation amplitudes, causing SIG-MATCH to fail. Since it turns out that the inputs to PF-19, the plan-fragment corresponding to the RF section, meet specifications, whereas its outputs do not, PF-19 becomes the CULPRIT. The localization process then enters a new level of recursion.

4.3 The loop problem solved.

This recursive localization process continues smoothly until arriving at PF-32, an instance of the plan, OSC. Figure 4.5 shows the various recursive activations of LOCAL. PF-32 got the blame because SIG-MATCH decided that the observed output of the oscillator did not match the expected output. (Note that the oscillator has no inputs apart from power.) In particular, the output is a quiescent DC bias. LOCAL tries the usual trick of expanding the OSC plan into its parts, finding the corresponding plan-fragments, and examining the i/o properties of those fragments. In the present situation, the tracing process finds itself in a loop. (Refer to figure 2.7.) This is because each part in the plan has a bad output <u>and</u> a bad input. In the explanation above of the tracing strategy, I pointed out that LOCAL could detect when it has closed a loop while tracing, but nothing was said as to how this state of affairs might be handled. Another specialist, LOOPS is invoked.

LOOPS[1] supplements LOCAL's hypothesis formation scheme based on tracing with other

----

1 The simple feed-back configuration of figure 4.6 presents a situation which might seem to need special attention by LOOPS. Suppose the feed-back control signal were intended to have an amplitude that varied inversely with the forward signal's amplitude. Suppose further that the control signal is large in spite of the large forward signal. Though this is a bug within a loop, it is evidently findable by LOCAL without recourse to LOOPS.

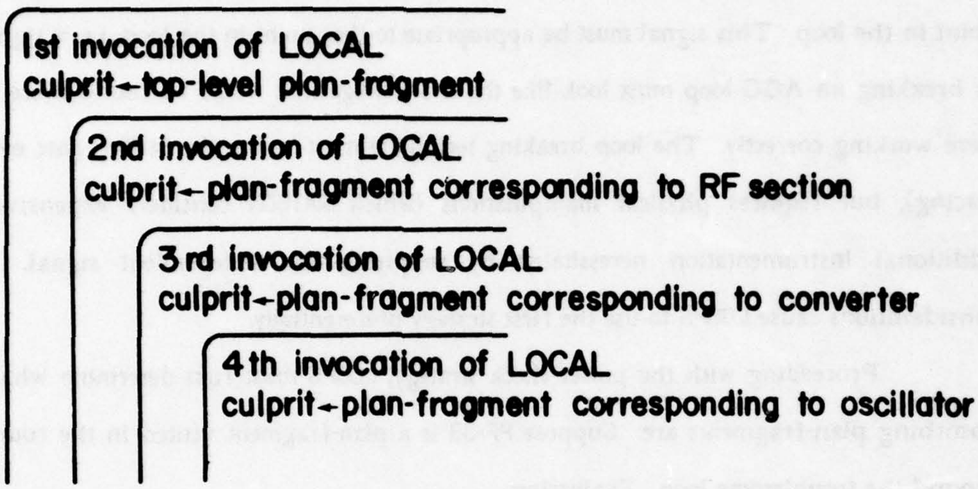

Figure 4.5 -- LOCAL's process state.

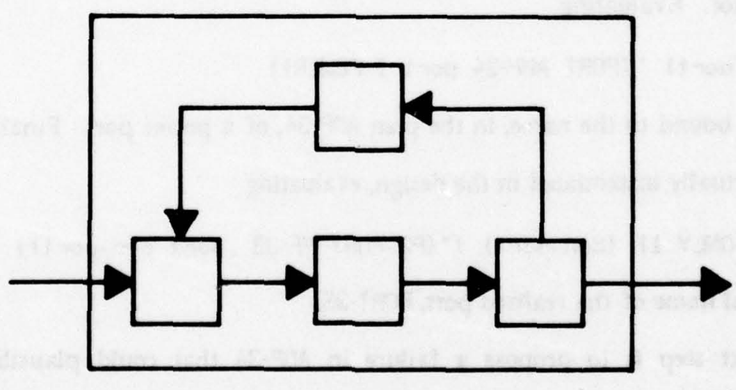

Figure 4.6 -- Three elements forward and one back.

schemes based on special knowledge about loop structures. LOOPS has two basic strategies. The first consists of looking at the power absorbing parts traversed in completing the loop and proposing failures based on knowledge of the purposes of those parts. The justification for this strategy is that powered parts have intrinsically higher *a priori* probabilities of failing because they encompass active components (and get hot), like transistors. The second strategy involves

breaking the loop by providing an externally (to the radio receiver) generated signal at some point in the loop. This signal must be appropriate to that point in the loop; i.e. a signal provided in breaking an AGC loop must look like the DC voltage that would be found there if the radio were working correctly. The loop breaking technique must isolate the failing part eventually (by tracing), but requires physical manipulations (which WATSON considers expensive), and the additional instrumentation necessitated by supplying the independent signal. The latter considerations cause LOOPS to use the first strategy preferentially.

Proceeding with the power-check strategy, LOOPS must first determine what the power absorbing plan-fragments are. Suppose PF-33 is a plan-fragment visited in the course of going around the troublesome loop. Evaluating

```
(FIND (THE-ONLY 1) (plan) '(PF PF-33 plan GR-78))
```

will yield the plan, AMP-34, of which PF-33 is an instance. AMP-34 is the plan of the amplifier used in the oscillator. Evaluating

```
(FIND ALL (port) '(PORT AMP-34 port ? POWER))
```

causes port to be bound to the name, in the plan AMP-34, of a power port. Finally, to determine how this port is actually instantiated in the design, evaluating

```
(FIND (THE-ONLY 1) (can-port) !"(PF-PORT PF-33 ,port can-port))
```

yields the canonical name of the realized port, PORT-35.

The next step is to propose a failure in AMP-34 that could plausibly explain the observed behavior of the oscillator. To do this LOOPS needs to know the *purpose* of the amplifier in the oscillator. This information can be obtained by evaluating

```
(FIND ALL (purpose) '(PF-PURPOSE (PF-32 PF-33 purpose))).
```

This returns a list of possible bindings of purpose, each element of which is a structure that looks like

(PF-PURPOSE PF-32 PF-33 predicate rule)

This says that the purpose of a particular part in a particular plan is to maintain the truth of some predicate via a computational rule. Specifically

```
(PF-PURPOSE PF-32 PF-33
   AMPLIFY-36
   (MAKE (> ~(AMP-GAIN PF-32) 1.)))
```

says that PF-33 in PF-32 serves to keep a certain parameter of PF-32 greater than 1. The rule used to maintain this condition is AMPLIFY-36. WATSON says to himself, "Suppose the predicate of the purpose were not true." The only way this could happen is for the maintaining rule to be deactivated. What would the effect on the oscillator be if AMPLIFY-36 were turned off?

To answer this question, LOOPS must pull out yet another structure having to do with teleology. It is an item of the form

```
(PF-GOAL
   plan-fragment rule list-of-parts-involved
   predicate list-of-results)
```

Before examining the various slots in the PF-GOAL item and considering a specific instance of such an item, let's step back for a moment and consider what we want to accomplish with this structure. Each abstract plan is made up of parts, each of the latter serving a purpose in the overall plan. Recall that plans are used to realize parts in yet larger plans. The parts at the next level have purposes too. Understanding a design presupposes understanding how purposes at one level of planning get mapped into purposes at the next lower level. The oscillator is a part in the converter plan (refer to figure 2.6). Its purpose there is to maintain an oscillatory signal of fixed frequency and amplitude. How does an oscillator, composed of amplifier, tank and amplitude stabilizer serve this purpose? A PF-GOAL item explains all.

A PF-GOAL may refer to a particular list of plan-fragments in cascade inside a distinguished plan-fragment. The goal of this cascade is to maintain the truth of a certain predicate. This predicate will be the same as the predicate in some PURPOSE item associated with

a plan-fragment at the next level up in the design. The resulting structure is a kind of trace, or scenario, of the operation of the parts of the plan. The result list together with the rules mentioned in the RESULTs listed are a kind of control structure that represents the flow of causality in the plan mentioned in the PF-GOAL item. This trace reveals how the predicate of the PF-GOAL is established by the interaction among parts. Let's turn again to the oscillator that fails to oscillate. It is associated with the item[1]

```
(PF-GOAL PF-32 RULE-37
  (PF-33 PF-38 PF-39)
  (STABILIZE
    ,~(AMPLITUDE CARRIER SIGNAL OSC-OUT PF-32)
    (AT ~(IDEAL-AMPL PF-32)))
  (RES-40 RES-41 ... ))
```

The sub-plan-fragment list (the fourth position in this item) specifies the plan-fragments on the loop that lead to the present invocation of LOOPS. The result list (the sixth position) reveals what would happen if the amplitude of the output signal were less than the desired amplitude, ~(IDEAL-AMPL PF-32), that is part of the design. The items

```
(RESULT RES-40
  (INIT)
  (< nowvalue ← ,~(AMPLITUDE CARRIER SIGNAL AMP-OUT PF-32)
     ~(IDEAL-AMPL PF-32)))
(RESULT RES-41)
  (INIT)
  ampl1 ← ,~(AMPLITUDE CARRIER SIGNAL OSC-GAIN PF-32))
(RESULT RES-42
  (INIT)
  ampl2 ← ,~(AMPLITUDE CARRIER SIGNAL OSC-OUT PF-32))
(RESULT RES-43
  (AMP-LIMITER-LAW-44 RES-40)
  (> ,~(AMPLITUDE CARRIER SIGNAL OSC-GAIN PF-32) ampl1))
(RESULT RES-45
  (AMPLIFY-36 RES-43)
  (> ,~(AMPLITUDE CARRIER SIGNAL OSC-OUT PF-32) ampl2))
(RESULT RES-46
  (TANK-LAW-47 RES-45)
  (> ,~(AMPLITUDE CARRIER SIGNAL AMP-OUT PF-32) nowvalue))
```

----

1 We shall eventually see that the structural relationship among PF-GOAL and PF-PURPOSE items is somewhat more complicated that the picture painted here. The items shown -- though not quite correct -- reveal the essential flavor of WATSON's analysis.

say that initially the tank input (refer to figure 2.7), tank output and limiter output take on values whose particulars are not important but they must be talked about, hence the Skolemizations ampl1 and ampl2. The filter output has an amplitude less than the desired value. This results in an increased limiter output, hence an increased amplifier output (filter input), and finally an increased filter output. WATSON incorporates an analyzer for such scenarios (see chapter 11) which is capable of detecting whether the prevailing rules (i.e. the ones active during the current activation of LOCAL) actually complete the scenario as as specified by the result list of the PF-GOAL item. Other PF-GOAL items account for how stability is achieved when the output amplitude of the oscillator is high or on target. LOOPS reasons (via scenario completion) that if the rule, AMPLIFY-36, were deactivated, the results depending on it would cease to be true. An initial result in the structure indicates that the proper amplitude has been undershot. A final result (depending on AMPLIFY-36) states that the amplitude has been increased[1], correcting for the initial state of affairs. Deactivating the rule, AMPLIFY-36, invalidates the result, leaving the amplitude in the initial state of undershooting the desired value, i.e. still at nowvalue. The result describing this latter state is matched against the actual state of affairs in the oscillator (Recall that it has zero output amplitude!) and is found to be consistent. LOOPS makes PF-33 the CULPRIT and calls LOCAL on it.

4.4 The indictment of Q301.

The application of LOCAL to PF-33 introduces another twist. Until now the localization process has not explicitly mentioned any of the electrical mechanisms that underly the functioning of a radio. The plans that have been dealt with thus far have all been of type CASCADE. This makes it possible to think about radios in terms of sequences of abstract signal processors. Of course, MAXWELL knows that the magnitudes of certain voltages and currents (as obtained from

----

[1] An initial result is a given; that is, it is dependent on the INIT rule and no other result. A final result is one on which no other result depends.

various instruments) correspond to various abstract signal descriptors. Still, electricity has been pretty far removed from the fault-finding process. AMP-34 is a plan of type CIRCUIT. A CIRCUIT plan makes explicit the voltage/current interactions among its parts in the same way that a CASCADE plan makes explicit the signal interactions. Unfortunately, CIRCUIT plans do not have the quasi-causal[1] behavior of a CASCADE. Consequently it is unclear where to start doing either the forward or backward reasoning that is characteristic of tracing. There are, nonetheless, other powerful heuristics that may be applied. WATSON has already made use of the "unreliability" of plans that have powered parts. In this final phase of localization, he will again make use of that heuristic. He will also make use of the observation that the AC and DC behaviors of a CIRCUIT plan are largely decoupled and therefore may generally be analyzed separately.

LOCAL first asks if there are any power consuming plan-fragments in PF-33. Such a query produces Q301-48, a plan-fragment instantiating a plan of type COMPONENT[2]. Like other plans, COMPONENT plans describe a <u>use</u> (of an atomic electronic part). As such the COMPONENT plan for a bi-polar transistor used in a common emitter configuration is different from one used in a common collector configuration. Hence the two uses would be described by different COMPONENT plans.

As usual, the next step after producing a candidate is checking to see if the candidate's outputs are consistent with its inputs. This is the point at which DC and AC analyses separate. Almost any failure in an active component will lead to important changes in the prevailing biases. Another important observation is that when thinking about the DC properties of plan-fragments at the CIRCUIT plan level, "input" and "output" are not very meaningful, since any of the branch or node variables may be considered "independent." LOCAL will worry about DC consistency --

---

1 A CASCADE plan embodies the sequential signal processing metaphor. Apart from the inconvenience of an occasional loop, this view of causality in a radio underlies the very powerful debugging tool, back-tracing.

2 Note that the mapping from a design onto a circuit diagram is done by instantiating COMPONENT plans, giving the resulting plan-fragments names identifiable with components on the diagram.

consistency with node biases on the circuit diagram and/or the consistency of the measured DC conditions of a component with the components terminal voltage/current description. Evaluating

```
(TEST '(DC-CONSISTENT Q301-48))
```

will reveal whether or not Q301-48 is in good shape with respect to its DC surroundings. Curiously enough, PF-33 is responsible for setting Q301-48's bias conditions (as is indicated by attached purpose commentary). The use of the JFET-49 plan, of which Q301-48 is an instance, requires that certain bias prerequisites be met. PF-33 takes responsibility[1] for meeting the obligation. It turns out that Q301-48's biasing is dependent upon the mode of operation of the oscillator. In particular, if the oscillator is not oscillating, the transistor will be biased for class A operation. Evaluating

```
(TEST '(OSCILLATING PF-32)) ;PF-32 is the oscillator plan-fragment
```

during the course of the consistency check will inform LOCAL that the oscillator is not oscillating, hence class A biasing of Q301. Unfortunately, the biases on the circuit diagram correspond to measurements done with the oscillator oscillating. So comparing the actual state of the circuit with the diagram will not give a valid consistency check as it might if Q301 were being run as, say, a garden variety class A, common source amplifier.

All is not lost, however. In using the JFET-49 plan, the AMP-34 expects certain i/o behavior at the terminals of Q301. Q301 is an MPF105 JFET. As such Q301 is quite accurately characterized by a certain table (which may be interpolated) of voltage/current measurements made at its terminals. LOCAL ADDs to the data base RESULTs that characterize the prevailing DC conditions:

```
(RESULT RES-50
   (INIT)
   (DC-BV (~(GATE Q301-48) Q301-48 ~(SOURCE Q301-48)) 0.)).
(RESULT RES-51
   (INIT)
   (DC-BV (~(DRAIN Q301-48) Q301-48 ~(SOURCE Q301-48)) 9.)).
```

----

[1] In particular, PF-33 contains parts whose purpose is to achieve the correct bias conditions.

The ADDition of these facts triggers the firing of a number of rules. A KVL rule associated with Q301-48 computes the drain/gate branch voltage, and various rules associated with the nodes distinguished in PF-33 compute node voltages with respect to the declared ground, NODE-52. The products of these computations are RESULTs too. There is one more rule associated with Q301-48 that is of immediate importance, the one representing its voltage/current characteristics. On hearing the ADDition of the branch voltage consequents, this last rule, VIC-53, computes the branch current in the source/drain branch of Q301

```
(RESULT RES-54
   (VIC-53 RES-50 RES-51)
   (DC-BC (~(DRAIN Q301-48) Q301-48 ~(SOURCE Q301-48)) 9.)).
```

Finally, the KCL rule associated with the source node, ~(SOURCE Q301-48), and the Ohm's law rule associated the plan-fragment for the source resistor, R302-55, compute respectively

```
(RESULT RES-56
   (KCL-57 RES-54)
   (DC-BC (~(SOURCE Q301-48) R302-55 NODE-52) .007))
```

and (remembering that NODE-52 is the ground node)

```
(RESULT RES-58
   (OHM-59 RES-60)  ;RES-60 is the result indicating
                    ; the branch voltage across R302.
   (DC-BC (~(SOURCE Q301-48) R302-55 NODE-52) .02)),
```

both of which are claims about the quiescent branch current through R302. Consistency monitoring (see section II.2) notes that RES-56 and RES-58 are incompatible. LOCAL is unleashed on Q301-48 which results in the transistor's being pulled for inspection. It is discovered that it has an opened gate.

5 The Case of C422

In this chapter I shall show much of the detail of how the scenario of the shorted capacitor is actually carried out by WATSON. There are two features of importance introduced here. The first is a technique that allows the consideration of non-linear signal processing in a very local fashion. The second is a new style of plan(-fragment) (whose type is called COUPLING) which engenders an important localization heuristic.

5.1 Getting started.

The complaint about the GR-78 is lodged by evaluating

```
(COMPLAINT GR-78
  (INPUTS
    (~(OBS-SIGNAL PORT-1)
      (CARRIER-COMPONENTS
        ((MODULATION AM)
         (CARRIER-FREQ x)
         (AMPLITUDE z ϵ {⌊.10⌋ *}|{⌈.10⌉ *})
         (MODULATION-COMPONENTS
           ((MODULATION-FREQ y ϵ {*}))))))))
  (OUTPUTS
    (~(OBS-SIGNAL PORT-2)
      (DISTORTION HARMONIC)
      (CARRIER-COMPONENTS
        ((MODULATION AM)
         (CARRIER-FREQ 0.)
         (MODULATION-COMPONENTS
           ((MODULATION-FREQ w • z ← y | ϵ {~ y}))
            (MODULATION-AMPL v • z ← 0. | ϵ {> 0.}))))))
  (CONTROL-BINDINGS
    (TUNING x ϵ {*})))
```

There are a number of new syntactic features introduced in this COMPLAINT. Notice the

declaration of the input signal amplitude, z. This says that z is to be some amplitude in the bottom ten or top ten percent of its allowable range, that is, a weak or a strong signal respectively. The features of interest in the output signal are specified by the declarations of the variables u and v. In particular

```
...u • z ← y | ε (~ y)...
```

says that the choice of a value for u depends on the way the value is chosen for z. If z is chosen from the generator (L.10⌋ *) (i.e. weak signals), u is assigned the same value as y. On the other hand, if z is chosen from the generator (⌈.10⌉ *) (i.e. strong signals), then u is chosen from the generator that produces harmonics of y (i.e. harmonic distortion).

```
...v • z ← 0.| ε (> 0.)...
```

says that v is a zero or non-zero amplitude, depending on the choice of z.

The degrees of freedom that were noted in section 4.1 prevail here as well. An additional degree of freedom is introduced by the declaration of z. Recall that in section 4.1 I remarked that WATSON generally does not manipulate the amplitude setting of the signal generator. In this case he must do so since the COMPLAINT depends upon signal strength. In order to determine how to select this parameter for the validation process, WATSON needs to know something of the sensitivity properties of the GR-78. Looking into the data base, he finds

```
(PLAN RECEIVER CASCADE)
(PF PF-3 RECEIVER GR-78)
(PF-PARAM PF-3 SENSITIVITY REG-61)
(VALUE REG-61 .0000001)
(PF-PARAM PF-3 FRONT-END-OVERLOAD REG-62)
(VALUE REG-62 .00002)
```

These items give the upper and lower bounds for reasonable signals to be given the GR-78. WATSON takes the geometric means of .1 and 2.1 μV, and 18 and 20 μV respectively, yielding six possible set-points for the input signal amplitude.

These set-points are combined with the set-points developed in section 4.1 to generate

the configurations for the validation test, this time a loop with eighty-one inner rounds. The forward reasoning for validating the complaint proceeds essentially as before. The DETERMINER for the frequency of the modulation output component reveals that 1 kHz should be present. WATSON applies his frequency analyzer to the output port, PORT-2, of the GR-78 and duly notes the lack of a 1 kHz signal, corresponding to input signal amplitudes of 2 μV or less, and the presence of 1 kHz fundamental plus significant harmonic components for input signals of amplitude 10 μV or more. Having thus validated the complaint, a test set-up must be chosen for the localization process. WATSON is inclined to choose a test configuration from among those that reveal extraneous behaviors (the presence of harmonics in this case) rather than from those showing a lack of output[1]. So he chooses an input signal of 5 MHz carrier, 1 kHz modulation, and an amplitude of 12 μV.

5.2 Localizing to the AGC.

He first queries the data base to see if he knows a bug that matches this situation. As before, the answer is negative and the tracing process begins. Successive applications of LOCAL reveal bad outputs from PF-19, and from PF-63, the plan-fragments corresponding to the RF section and IF strip respectively. SIG-MATCH fails at the output port of the RF section because MAXWELL reveals that there are several harmonics of the 1 kHz modulation, whereas the rule for PF-19 demands that only the fundamental of the 1 kHz modulation should be present. SIG-MATCH also fails at the output port of PF-63 because the carrier frequency of the observed output does not match the expected frequency -- 455 kHz. Analysis of the output signal reveals a 455 kHz carrier component plus a 1 kHz modulation with various harmonics of that modulation. Moreover, there are harmonics of the 455 kHz carrier. The audio frequency harmonics are not present in the input signal from the converter. Therefore the observed output does not match the

1 There are likely to be many underlying causes for a lack of output, whereas, spurious output behavior may offer guidance as to the source of the wayward output.

expected output which has only the 455 kHz carrier and the 1 kHz modulation. The question then becomes whether or not the inputs to the IF strip are correct.

Proceeding to check the inputs of the IF strip, its main input, from the converter (see figure 2.5), is judged correct via forward reasoning facilitated by determiners. In attempting to match the auxiliary control input coming from AGC2 through port PORT-64 against the expected signal at that port, the control bias that is supposed to be developed there falls short of the expected value. This discovery triggers a *caveat alarm*[1] that informs LOCAL of the incorrect use of the plan, a token of which the IF strip (a part) is bound to. Usually the determination of the correctness of the output of a plan-fragment, given its inputs, is a straightforward use of determiners, rules, and signal back-tracing. In this case it is complicated by the fact that one of the inputs to the PF-63 directly contradicts a caveat, thereby failing to meet conditions imposed for the correct use of PF-63's type. Hence it cannot be guaranteed that the rules that usually describe the plan's input-output behavior remain valid. In order to find out what should happen in these adverse circumstances, WATSON must look inside PF-63.

The first order of business is to discover the source of the caveat, i.e. for what object in the design is the condition of the caveat being demanded? This is easily determined. When the caveat alarm went off, the variable cave-canem was bound to the name of the offended caveat. Evaluating

```
(FIND (THE 1) (pred dem) !"(CAVEAT ,cave-canem pred dem))
```

reveals the condition the caveat wanted to be true, i.e. the value of pred, and its source, i.e. the value of dem. dem is bound to PF-65, the plan-fragment which realizes the part of the IF strip called IF1 (see figure 2.10), an amplifier. In order to determine what the IF strip does when the condition of the caveat is not met, WATSON needs an explanation of the caveat. What is the nature of such an explanation?

- - - - - - - - - - -

[1] The triggering is done via IF-ADDED methods. The declaration of a caveat in the assembly of a design (see chapter 10) engenders the creation of a collection of IF-ADDEDs that watch over the maintenance of the condition demanded by the caveat. When a contrary condition is asserted, LOCAL is interrupted and informed of the trouble.

In order to approach this last question properly, we need to understand better where the caveat came from in the first place and how it came to be visible immediatedly inside the plan for the RF section. PF-65, the plan-fragment corresponding to IF1, is an instance of an amplifier used in a class A regime. The type of this plan-fragment claims to have an externally supplied bias used to control the gain of the amplifier. The input-output design specifications associated with this type say that for control biases greater than 1.0 V it is a class A amplifier, for biases between 0.6 and 1.0 V it is a class B amplifier, and finally biases less than 0.6 V yield a class C amplifier. Consequently it is the use of this plan in the IF strip that is the source of the caveat. The caveat's becoming visible external to the IF strip plan-fragment, PF-63, inside the RF section is a result of the design assembly process. (See chapter 10 for further details.) PF-65 has a prerequisite called a *requirement*, which describes the bias to be supplied at PF-65's gain-control port that maintains its "class-A-ness." A requirement is similar to another class of prerequisites called *needs*. The former must be be satisfied at "run-time" while the latter usually must be satisfied at design-assembly-time. In either case the satisfaction is guaranteed by the compilation and activation of IF-ADDED methods. In the present case, when MAXWELL ADDs to the data base the value of the bias measured at the port between the AGC2 and the IF strip, the method runs, reporting the problem to LOCAL. Note that though this method is compiled in the course of instantiating the plan of IF1 (yielding PF-65), this method in fact "listens" at the level of the IF strip. This is a consequence of the fact that the external bias port of the IF1 amplifier is also an external bias port of the IF strip. Hence, if there is trouble, it will be recognized at the first opportunity.

An explanation of the caveat, then, should allow WATSON to infer that IF1 is actually operating in a class C regime rather than the intended class A regime. WATSON first asks which of the possible uses of AMP2, the type of PF-65, matches the prevailing quiescent conditions surrounding PF-65 in the GR-78 circuit. The evaluation of

```
(OP-SPEC-MATCH 'PF-65 cave-canem)
```

yields a list of methods that give an imperative description of IFI. These methods embody the rule -- called AMPLIFY-66 -- and the determiners needed for predicting its input/output behavior under the prevailing circumstances. OP-SPEC-MATCH expects to find that the plan of which PF-65 is an instance has a number of alternative modes of operation. The caveat reflects having chosen one of those modes for use in the GR-78 design. The details of how OP-SPEC-MATCH does its job are tangential to WATSON's present line of thinking, so they will be postponed until section 10.7.

WATSON now has in his hands the behavioral description for IFI that actually applies -- in contrast to the description called for in the design. What does he do with it? He arrived at this point because he was attempting to rationalize the output of the IF strip in terms of its inputs. Ordinarily (apart from having set off the caveat alarm) he would simply have ground away with the necessary forward reasoning. Unfortunately, the rule for the IF strip relies on the rule for IFI. But WATSON has a new rule for IFI, included in the list of methods returned by OP-SPEC-MATCH. The IF strip, itself a CASCADE plan, is part of a CASCADE plan. Conceptually WATSON simply forgets (temporarily) the separate identity of the IF strip and inserts in its place the cascade inside the IF strip. The ability to do this transformation and use it to make sensible predictions relies on the fact that the signal descriptions generated by the rules for the IF strip's parts are compatible with the descriptions generated by the rule for the IF strip itself. Having thus *flattened* the IF strip modulo the IFI amplifier (a detailed account of which is given in section 11.4), WATSON carries out the following recipe:

1. Disable all the methods associated with PF-63 -- the plan-fragment for the IF strip.
2. Enable all the methods associated with the parts of the IF strip[1].
3. Garbage collect all the items deduced using determiners that no longer apply.
4. Now do the forward reasoning to predict the behavior of the "phantom" IF strip.

----

[1] These would normally be activated when LOCAL was applied to PF-63. Turning them on inside the RF section is the essence of inserting the cascade composing the IF strip in its place.

Point '3' is technically complex, but not difficult, and depends on the use of the context machinery together with annotations on data base items indicating how they were formed in the first place.

Now WATSON can finish off the job of inferring that the mis-bias at the control port of the IF strip can cause the problem. At the output port of the IF1 amplifier the forward reasoning just done predicts the presence of the first and second harmonics of the carrier, 455 kHz, each having a superimposed modulation consisting of the first five harmonics of 1 kHz[1]. The signal is propagated through the rest of the IF strip via the usual kinds of "linear" rules that we have seen before. SIG-MATCHing the predicted output of the IF strip against the output observed at the IF strip's output port completes a successful localization at this level.

5.3 Indicting C422.

At this point WATSON is pretty convinced that the source of the difficulty lies inside AGC2. Let's look into the circuit underlying AGC2 as illustrated in figure 5.1. It is a class B amplifier followed by an RC coupling network. The purpose of the circuit is to compute a time-averaged signal strength. The period with respect to which the time averaging is done is determined by the RC network of figure 2.12 and is about .1 seconds in this case. The mechanism by which this time averaging is achieved is not unlike demodulation. The class B amplifier rectifies the modulated signal, and the RC network does peak detection on the result, giving the largest value seen in the last .1 seconds. The class B amplifier supplies two other important functions as well. It is a buffer amplifier, thereby preventing excessive loading of the input to the detector of the GR-78. Also, being a common emitter configuration, it is inverting. Hence the

---

1 The spectral analysis embodied in the AMPLIFY-66 rule is arbitrarily limited to thinking about the first two harmonics of the carrier and the first five harmonics of the modulation. This restriction is imposed because the presence of large harmonics components of this variety are sufficient evidence to indicate that observed outputs are explained solely by the inputs -- i.e. there is no problem inside IF1.).

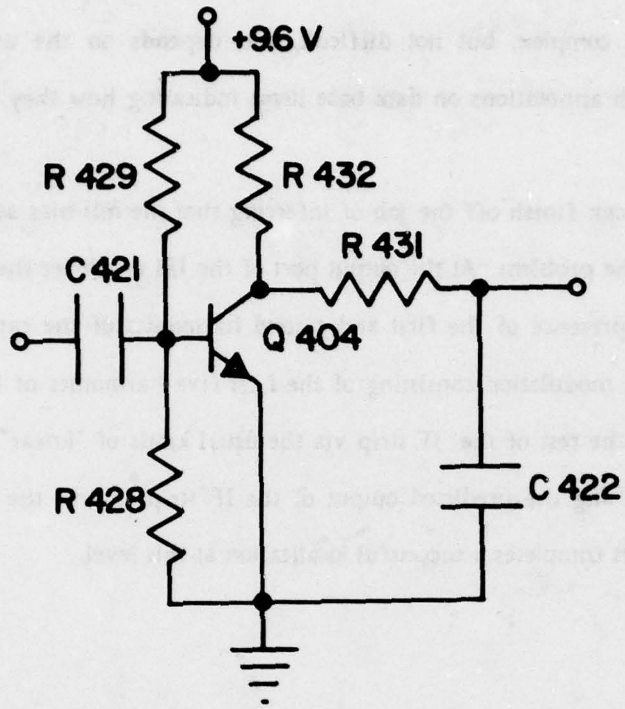

Figure 5.1 -- The circuit details of AGC2.

magnitude of the amplifier's output varies inversely with the size of its input. This is exactly what is wanted since the whole point of the AGC feed-back loop, of which AGC2 is a part, is to stabilize the strength of the signal delivered at the output of the IF strip in the face of any drifting in the signal strength at the input.

As I pointed out in section 1.4 the viability of tracing as a localization strategy relies upon the unilateral nature of the typical plan(-fragment). The RC network of figure 2.12 is not unilateral. In particular the voltage from A to ground affects and is affected by the voltage from B to ground. To embody this notion, among others, WATSON is aware of a class of plans I have not previously mentioned -- COUPLING plans. The essential feature of such plans is that they modify the relative concentrations of the various spectral components in their input signals.

Voltage measurements across either the input or output port[1] of such a plan would reveal identical signals modulo some DC offset. This latter feature suggests that it meaningless to ask whether the output "follows" the input since they are not causally related. We will see shortly how this apparent bug in tracing can be used to WATSON's advantage.

To return to WATSON'S analysis, he sees the CASCADE plan of figure 2.11 and begins the trace as usual but notices that the first part on the signal path is bound to a COUPLING plan-fragment. Instead of verifying that the output from the RC network is that which would be determined by its inputs, he checks to see if the output is the same as the inputs modulo the DC offset associated with the plan-fragment. This can be done simply by SIG-MATCHING the two observed signals. The match fails because the output of the filter is a DC voltage whereas the input is a rectified, amplified and inverted copy of the input to AMP. So LO-PASS is not meeting its specifications.

LOCAL is applied to the plan-fragment, PF-67, to which LO-PASS is bound. From the point of view of hypothesis formation and forward reasoning PF-67 is a token of a CIRCUIT plan. Since this plan contains no powered parts (a general feature of COUPLING plans), DC analysis is not particularly interesting, hence only incremental analyses are carried out. Contemplating what part to propose as failing, LOCAL makes a choice based on *a priori* probabilities of failure. In particular capacitors are likely to fail before resistors. Further C422 is more likely to short than to open. WATSON hypothesizes the shorting of C422 and imagines what would happen if an incremental voltage increase were applied at A of figure 2.12 by ADDing

```
(RESULT RES-68
  (INIT)
  (AC-BV
    (~(A PF-67) ~(RES PF-67) ~(CAP PF-67)  ~(GND PF-67))
    (UP (TOWARD 9.6))))).
```

The Ohm's law rule for R431 runs and ADDs

----

[1] The ports of such plans are generally voltage ports. The unwanted portion of the input signal is shunted into the local ground, letting the "good stuff" through.

```
(RESULT RES-69
   (OHM-70 RES-68)
   (AC-NV ~(B PF-67) (UP (TOWARD 9.6)))).
```

If the capacitor C422 were doing its intended job, it would have frustrated R431's attempt to pull up node B by hiding RES-69 and asserting

```
(RESULT RES-71
   (DV/DT-72 RES-69)
   (AC-NV ~(B PF-67) UNCHANGED))
```

A similar analysis results if the voltage across the input port of LO-PASS were to go down incrementally.

WATSON's hypothesizing the shorting of C422 results in the disabling of the old rule, DV/DT-72, and the enabling of a new rule, CAP-BUG-RULE-73. Part of the enabling of this rule is the declaration of node B as an incremental ground. So if the same initial result as above were asserted in the context of the new rule, OHM-70 would assert that the current through R431 increases, but the node voltage at B is unchanged, which is precisely the symptom observed. C422 is pulled from the circuit and is discovered to be shorted.

6 A Case of Misalignment

In this chapter the details of the debugging of the misaligned front-end are revealed. In particular, the data structures underlying the concept of alignment are explained and WATSON's alignment expert, ALIGN, makes its debut. As with WATSON's other expert, LOOPS, ALIGN is brought to bear because of local plan structure and behavior.

6.1 The usual preliminaries.

WATSON is introduced to this new problem by telling him

```
(COMPLAINT GR-78
  (INPUTS
    (~(OBS-SIGNAL PORT-1)
      (CARRIER-COMPONENTS
        ((MODULATION AM)
         (CARRIER-FREQ x ← *SOME)
         (MODULATION-COMPONENTS
           ((MODULATION-FREQ y ϵ {*})))))))
  (OUTPUTS
    (~(OBS-SIGNAL PORT-2)
      (CARRIER-COMPONENTS
        ((MODULATION AM)
         (CARRIER-FREQ 0.)
         (MODULATION-COMPONENTS
           ((MODULATION-FREQ y)
            (MODULATION-AMPL z ϵ {> 0.}))))))))
  (CONTROL-BINDINGS
    (TUNING w ← *SOME ≠ x)))
```

The variable x is existentially specified as some particular carrier frequency. Any valid modulation frequency is assignable to y. The key feature to notice is that when the radio is tuned to some frequency w other than the broadcast frequency, x, the radio puts out a detectable audio

signal at the modulation frequency. y, as usual is assigned 1 kHz. z, being an output variable, will match any measurable amplitude. The assignments of x and w -- they being existentially specified -- are somewhat problematic. In principle, WATSON might have to examine all pairs of assignments to x and w in order to come up with an example that validates the COMPLAINT. In practice he takes the optimistic view that the existential specification indicates that validating values are easy to find. He arbitrarily sets x to 5 MHz and does a sweeping search downward for an appropriate value for w. The search is stepped at 100 kHz intervals, search stepping intervals and limits being chosen for convenience, keeping in mind the GR-78's receiving spectrum. At each of the receiver settings (the volume control being at its usual middling setting) WATSON carries out the usual SIG-MATCHes, and is successful at every tuning until 4.1 MHz, at which point he finds a significant 1 kHz component at the output. In particular with the generator broadcasting at 5.0 MHz, the receiver tuned to 4.1 MHz, SIG-MATCH fails because there is observed a 1 kHz output when there should be none at all. Since he is in search mode, WATSON tries to optimize the find by making incremental variations in frequency around 4.1 MHz. This is done in increments of 10 kHz (because of channel width considerations imposed by civil law) and the amplitude of the 1 kHz output is maximized at 4.09 MHz. This configuration is the test set-up to be used in debugging.

6.2 Localizing the problem.

The back-trace proceeds as usual. SIG-MATCHing expected and observed signals reveals bad output from and input to the AF section (see figure 2.4). CULPRIT is therefore bound to PF-19, the plan-fragment for the RF section, and LOCAL is applied to it. Observation at the output of the detector (see figure 2.5) shows a 1 kHz modulation, which does not match the predicted output. But the input to the detector is a 455 kHz carrier modulated by 1 kHz, which by the rule for the detector transformation implies that the detector's output is reasonable. Similar reasoning applies to the IF strip. Its output, though wrong in principle, agrees with its input.

A caveat alarm is triggered because in the course of getting the observed input signal for the converter -- the plan-fragment for which is PF-74 -- from MAXWELL, a signal structure is generated having a carrier of 5.91 MHz. The alarm mechanism, as in the previous scenario, is the product of design teleology. As before, the sounding of this alarm binds cave-canem to the offended caveat and allows WATSON to find the demander of the caveat by evaluating

```
(FIND (THE 1) (pred dem) !"(CAVEAT ,cave-canem pred dem))
```

which binds dem to PF-75, a token of the plan, MX, the plan for the mixer inside the converter. Evaluating

```
(OP-SPEC-MATCH 'PF-75 cave-canem)
```

returns a list containing the methods that actually govern the mixer's behavior in the face of images. As in the case of the IF strip, the converter is flattened (modulo the mixer) to account for its behavior in the presence of an image station. The recipe of section 4.3 again applies. Reasserting the observed sign descriptors at the phantom converter's *input* reveals that the image station should come through, exactly as observed. WATSON's attention turns to the apparent source of the 5.91 MHz signal, the RF amplifier.

The input to the RF amplifier is known to be correct since it comes from the signal generator. The structure underlying PF-76 is illustrated in figure 2.13. As usual, WATSON's inclination is to back-trace, but when MAXWELL is asked for the signal data, WATSON receives a polite refusal of the form

```
(CANT-MEASURE portname BECAUSE reason).
```

In particular the *reason* substructure (for the output filter) would be

```
(CHANGE ~(BAND-PASS PF-77)) ;PF-77 is the plan fragment of the output filter
```

indicating that attempting to do measurements at the ports of the output filter changes its band-

pass characteristic[1]. It should also be pointed out, however, that even if the measurements expert had not complained, the plans associated with the input and output filters are of type COUPLING, hence not likely to be helpful in direct causal analysis.

Fortunately WATSON already noticed something that will *prove* helpful in the present circumstances. Recall that when he was looking into the cause of the caveat alarm, pred was bound to an explanation of why the alarm was sounded. The explanation is[2]

```
(FALSE
 (CARRIER-REJECTION
  (HIGH-BAND <<<.,~(FREQUENCY CARRIER SIGNAL OSC-IN PF-75)
                  - .,~(FREQUENCY CARRIER SIGNAL RF-IN PF-75)> 2>
                + <.,~(FREQUENCY CARRIER SIGNAL RF-IN PF-75)>>)
      (AT RF-IN)));
```

meaning that PF-75 expected no image signal to be at the port between the RF amplifier and the converter. The computation of the image carrier frequency embedded in this explanation of the caveat's cause for alarm yields a value of 5.91 MHz. WATSON asks the obvious question. Does PF-19 have anything to do with meeting the requirement implied by the caveat? Evaluating

```
(FIND ALL (rule part)
  '(PF-PURPOSE  PF-19 part rule
     (CARRIER-REJECTION
      (HIGH-BAND <<<.,~(FREQUENCY CARRIER SIGNAL OSC-IN PF-19)
                     -.,~(FREQUENCY CARRIER SIGNAL RF-IN PF-19)> 2>
                   + <.,~(FREQUENCY CARRIER SIGNAL RF-IN PF-19)>>)
      (AT RF-IN)))
```

yields an affirmative answer -- the RF amplifier plan-fragment (the value of part). Another FIND can be done to determine what parts in the RF amplifier support this goal. WATSON discovers that the input and output filters of the of the RF amplifier are active participants in the achievement of the goal.

- - - - - - - - - - -

1 This message, incidentally, is a result of WATSON's reading commentary on PF-77 concerning its input impedance. As we shall later see, this is the same kind of commentary as is used by LOOPS when breaking loops.

2 A pathname construct preceded by '.' is equivalent to a path name whose first path identifier is VALUE, i.e. .~(FREQUENCY ...) is equivalent to ~(VALUE FREQUENCY ...). Also, the bracketting characters, '<' and '>', are introduced to delimit subexpressions of the descriptive formulae used by WATSON. Since much of his descriptive notation is infix in nature, parenthesization distinct from that used for list structure is necessary.